



US007660828B2

(12) **United States Patent**  
**Johnston et al.**

(10) **Patent No.:** **US 7,660,828 B2**  
(45) **Date of Patent:** **\*Feb. 9, 2010**

(54) **SYSTEMS AND METHODS FOR GENERATING MARKUP-LANGUAGE BASED EXPRESSIONS FROM MULTI-MODAL AND UNIMODAL INPUTS**

5,884,249 A 3/1999 Namba et al.  
6,073,098 A 6/2000 Buchsbaum et al.  
6,233,544 B1 5/2001 Alshawi  
6,324,512 B1 11/2001 Junqua et al.  
6,401,065 B1 6/2002 Kanevsky et al.  
6,807,529 B2 10/2004 Johnson et al.  
6,823,308 B2 11/2004 Keiller et al.  
6,868,383 B1 3/2005 Bangalore et al.

(75) Inventors: **Michael J. Johnston**, Hoboken, NJ (US); **Srinivas Bangalore**, Morristown, NJ (US)

(Continued)

(73) Assignee: **AT&T Intellectual Property II, LP.**, New York, NY (US)

**FOREIGN PATENT DOCUMENTS**

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 297 days.

EP 0 243 009 10/1987

(Continued)

This patent is subject to a terminal disclaimer.

**OTHER PUBLICATIONS**

Johnston et al., "MATCH: multimodal access to city help", retrieved from <http://www.research.att.com/~johnston/johnstonelatasru01.pdf> on Feb. 17, 2006, published on Oct. 4, 2001, 4 pages.

(Continued)

(21) Appl. No.: **11/773,659**

(22) Filed: **Jul. 5, 2007**

*Primary Examiner*—Charles Rones  
*Assistant Examiner*—Yuk Ting Choi

(65) **Prior Publication Data**

US 2008/0046418 A1 Feb. 21, 2008

(57) **ABSTRACT**

**Related U.S. Application Data**

(63) Continuation of application No. 10/278,921, filed on Oct. 24, 2002, now Pat. No. 7,257,575.

When using finite-state devices to perform various functions, it is beneficial to use finite state devices representing regular grammars with terminals having markup-language-based semantics. By using markup-language-based symbols in the finite state devices, it is possible to generate valid markup-language expressions by concatenating the symbols representing the result of the performed function. The markup-language expression can be used by other applications and/or devices. Finite-state devices are used to convert strings of words and gestures into valid markup-language, for example, XML, expressions that can be used, for example, to provide an application program interface to underlying system applications.

(51) **Int. Cl.**  
**G06F 17/30** (2006.01)

(52) **U.S. Cl.** ..... **707/104.1**

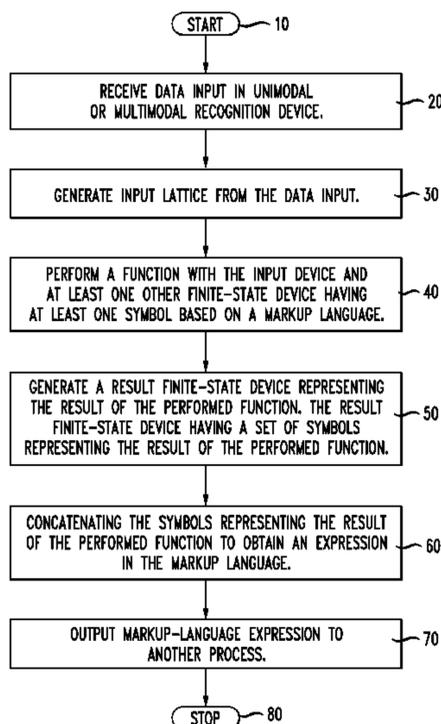
(58) **Field of Classification Search** ..... None  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

5,748,974 A 5/1998 Johnson  
5,781,179 A 7/1998 Nakajima et al.

**24 Claims, 20 Drawing Sheets**



U.S. PATENT DOCUMENTS

6,961,776	B1 *	11/2005	Buckingham et al. ....	709/229
7,107,205	B2	9/2006	Kempe	
2002/0161907	A1 *	10/2002	Moon .....	709/230
2003/0046087	A1	3/2003	Johnston et al.	
2003/0055644	A1	3/2003	Johnston et al.	
2003/0065505	A1	4/2003	Johnston et al.	
2003/0093419	A1	5/2003	Bangalore et al.	

FOREIGN PATENT DOCUMENTS

GB	2 209 418	5/1989
----	-----------	--------

OTHER PUBLICATIONS

Wu et al., "Statistical multimodal integration for intelligent HCI", Neural Networks for Signal Processing IX, 1999. Proceedings of the 1999 IEEE Signal Processing Society Workshop, Madison, WI, USA Aug. 23-25, 1999, Piscataway, NJ, USA, IEEE, US, Aug. 23, 1999, pp. 487-496, XP010348491 ISBN: 0-7803-5673-X \* p. 489, line 27—p. 490, line 27.

Mohri et al., "Weighted Automata in Text and Speech Processing" Extended Finite States Models of Languages: Proceedings of the ECAI 96 Workshop, [Online] 1996, XP002377632 Retrieved from the Internet: URL:<http://citeseer.ist.psu.edu/mohri96weighted.html>>[retrieved on Apr. 20, 2006] p. 1, paragraph 2—p. 2, p. 3, paragraph 4—p. 4.

Johnston et al., "Finite-state Methods for Multimodal Parsing and Integration" in ESSLLI Workshop on Finite-State Methods, Helsinki, Finland, Aug. 2001, pp. 1-6.

Mohri, "Finite-State Transducers in Language and Speech Processing", 1997 Association for Computational Linguistics, vol. 23, No. 2, pp. 1-42.

Johnston, "Multimodal Language Processing", Proceedings of ICSLP 1998 (published on CD-ROM only).

Johnston et al., "Finite-state Multimodal Parsing and Understanding", In Proceedings of COLING 2000, Saarbruecken, Germany.

Johnston, "Unification-based Multimodal Parsing", In Proceedings of COLING-ACL, 1998, pp. 624-630.

Johnston, "Unification-based Multimodal Integration", Proceedings of the 35<sup>th</sup> ACL, Madrid, Spain, 1997, pp. 281-288.

Pereira et al., "Finite-State Approximation of Phrase Structure Grammars", in E. Roche and Schabes Y., editors, Finite State Devices for Natural Language Processing. MIT Press, Cambridge, Massachusetts, pp. 246-255.

Johnston, "Deixis and Conjunction in Multimodal Systems", Proceedings of COLING 2000, Saarbruecken, Germany, 2000.

Co-pending U.S. Appl. No. 10/278,921 to Johnston et al., filed Oct. 24, 2002, entitled "Systems and Methods for Generating Markup-Language Based Expressions from Multi-Modal and Unimodal Inputs."

\* cited by examiner

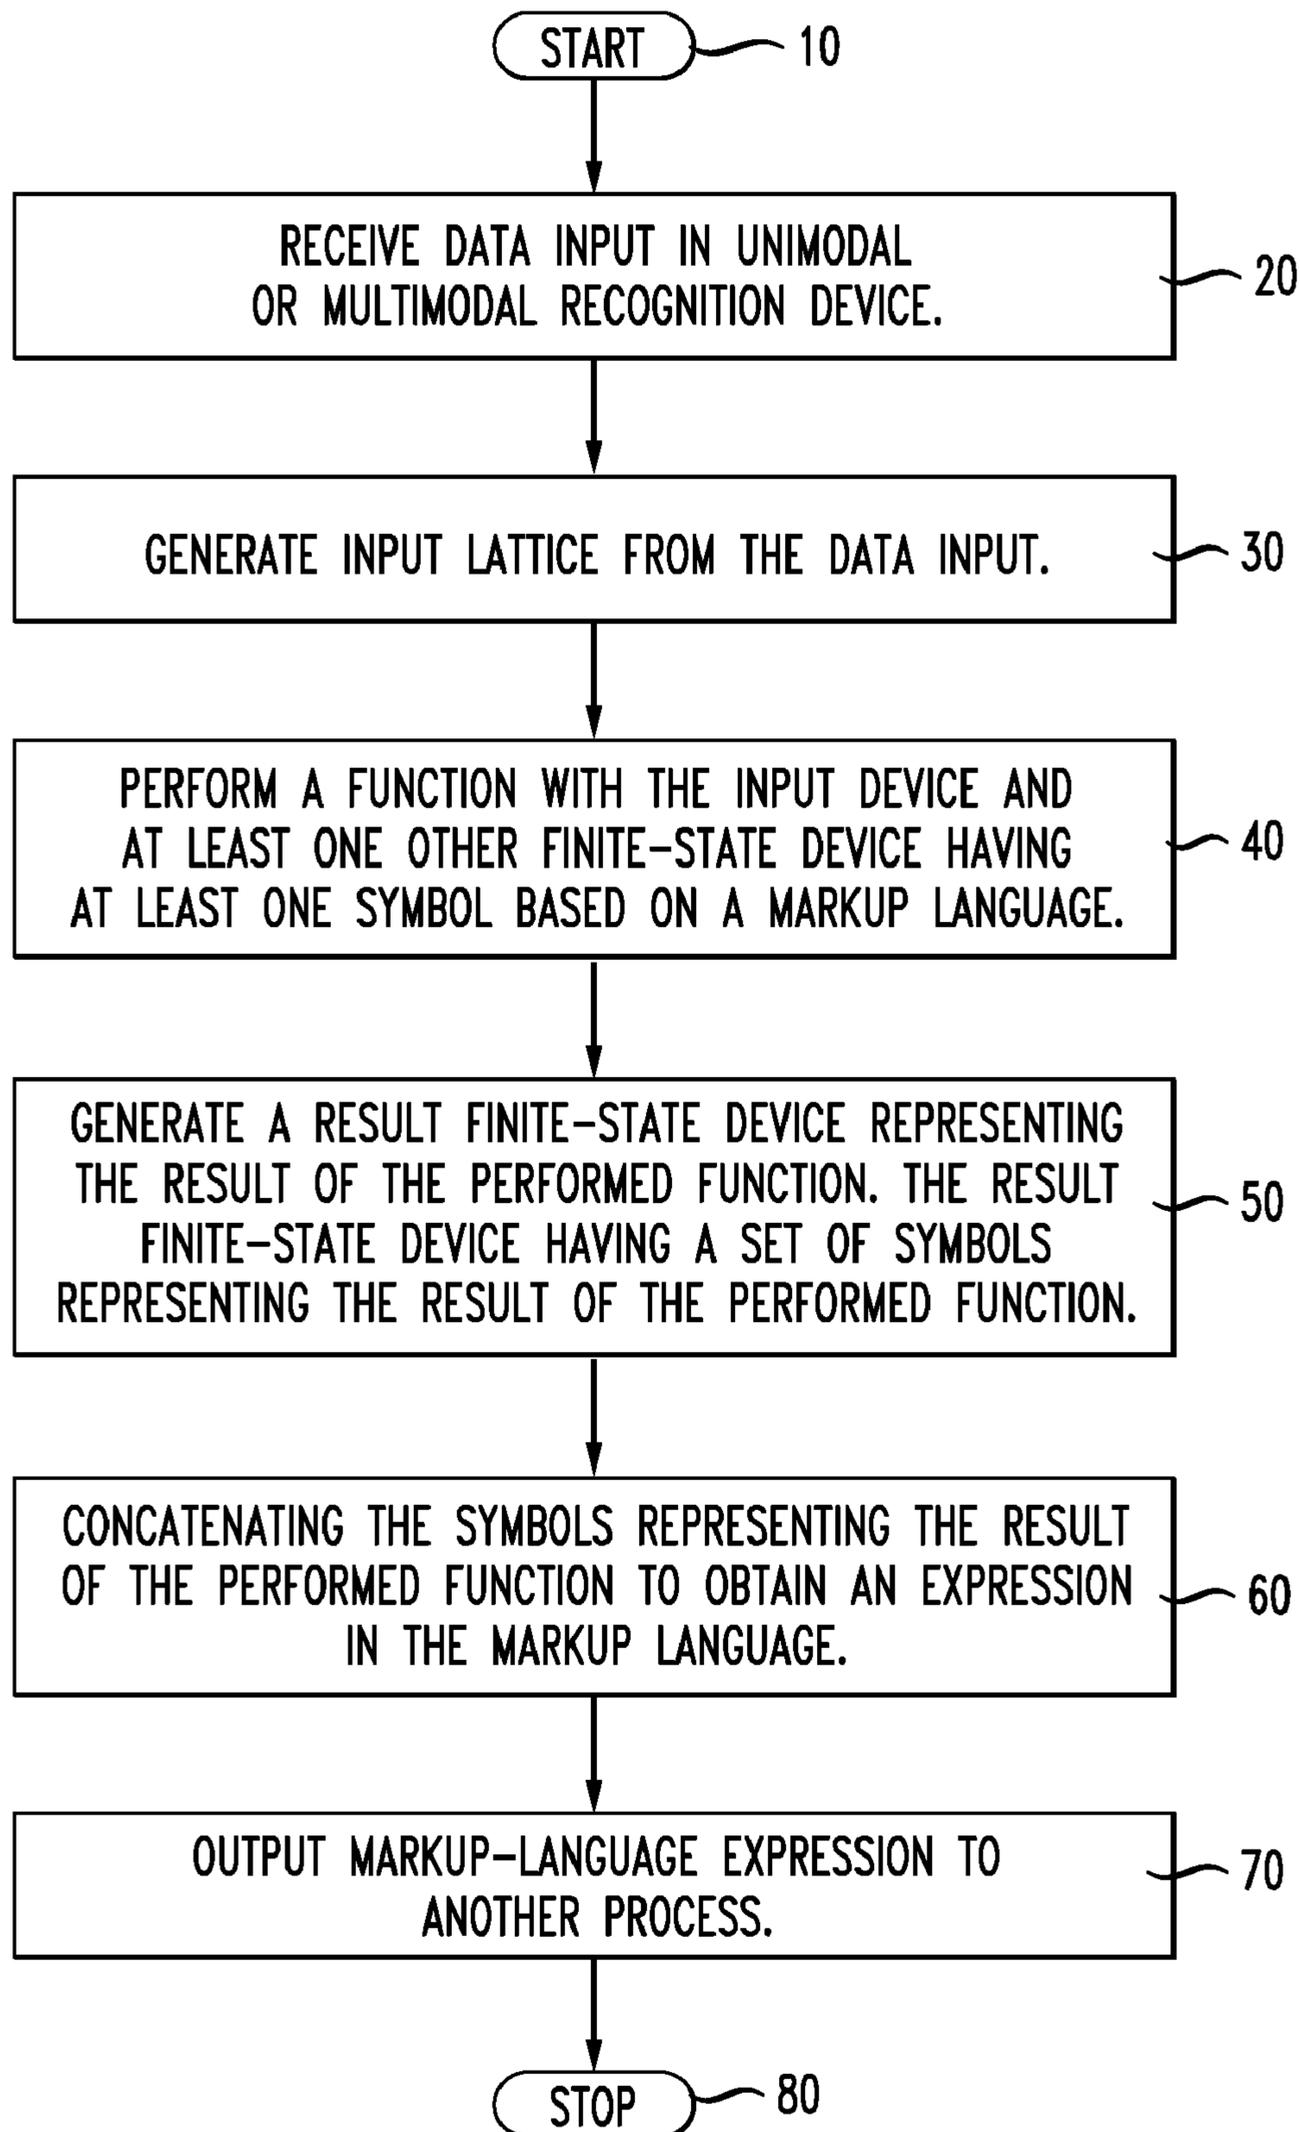
*FIG. 1*

FIG. 2

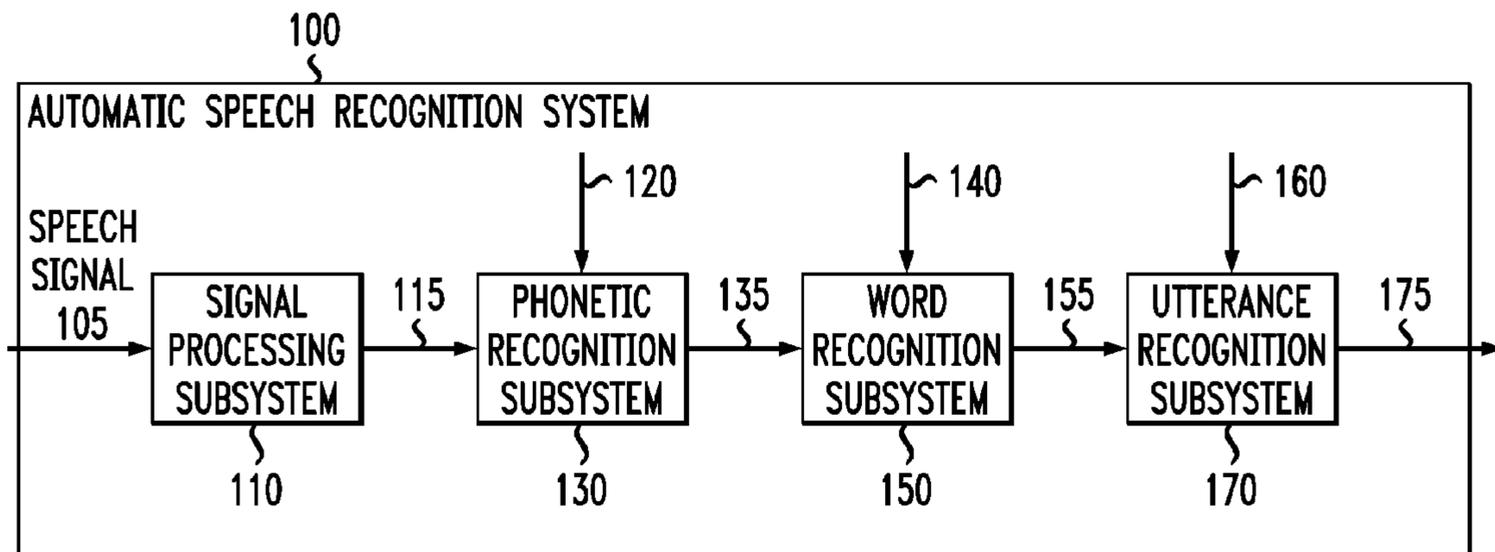


FIG. 3

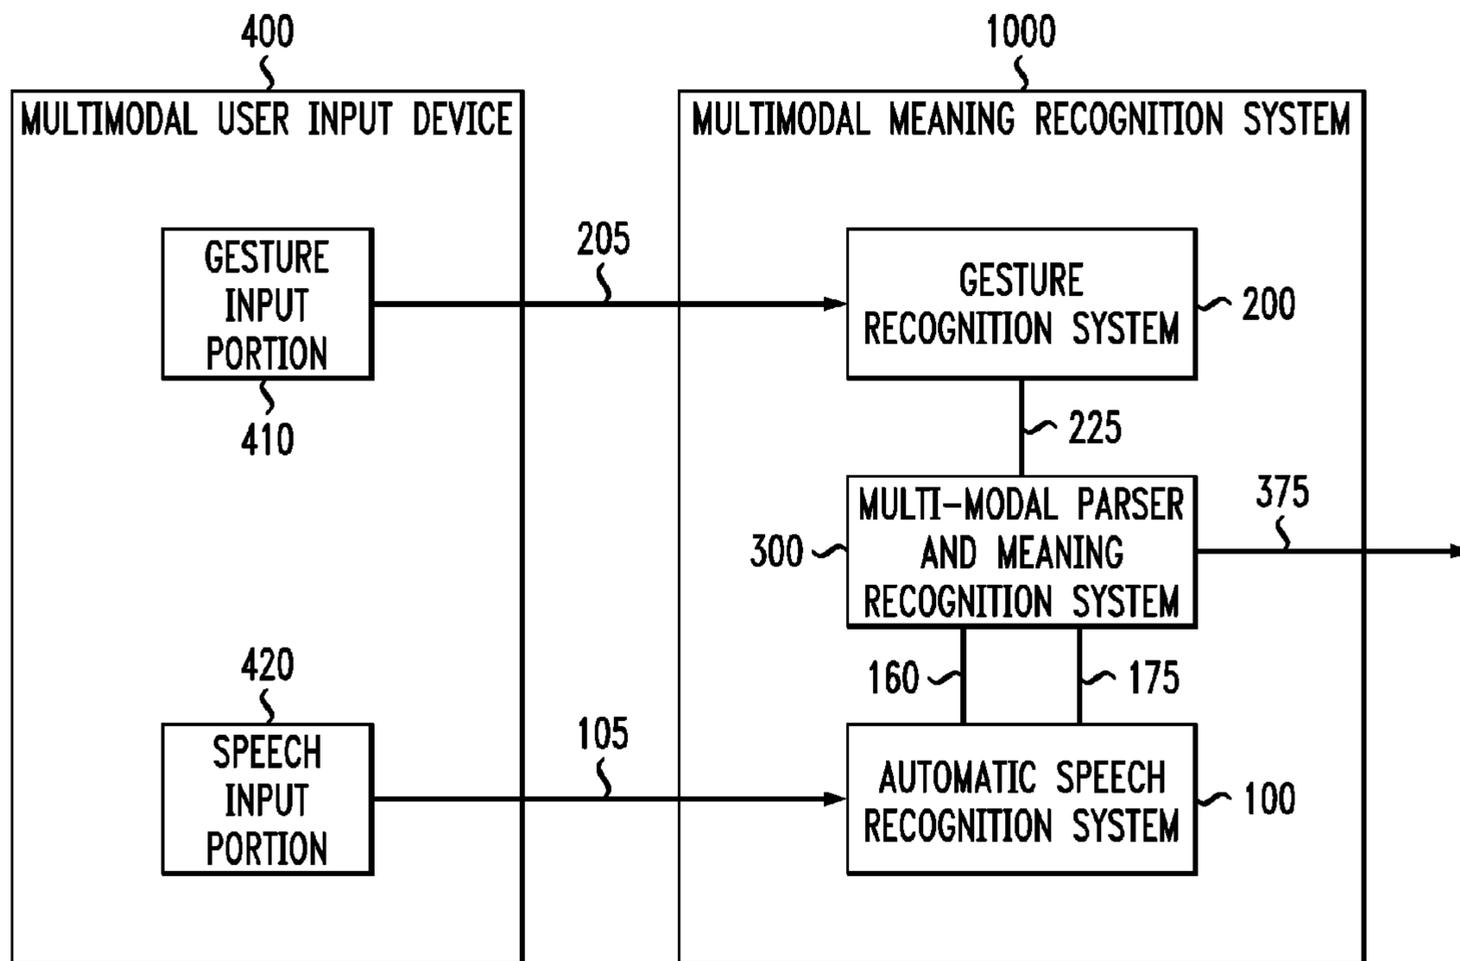


FIG. 4

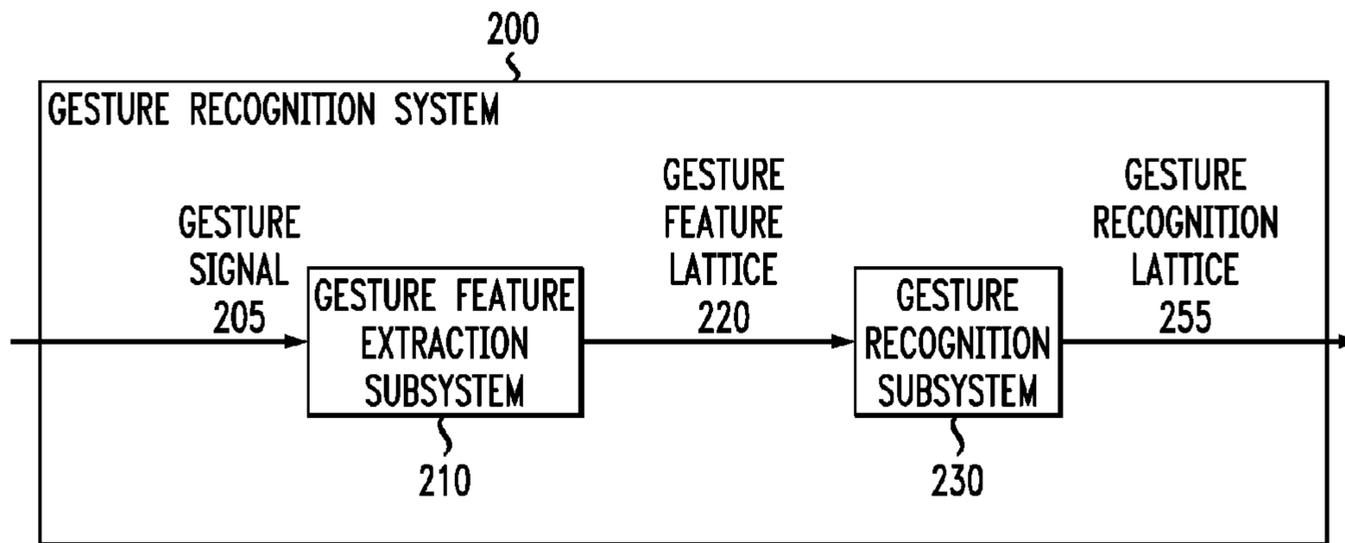


FIG. 5

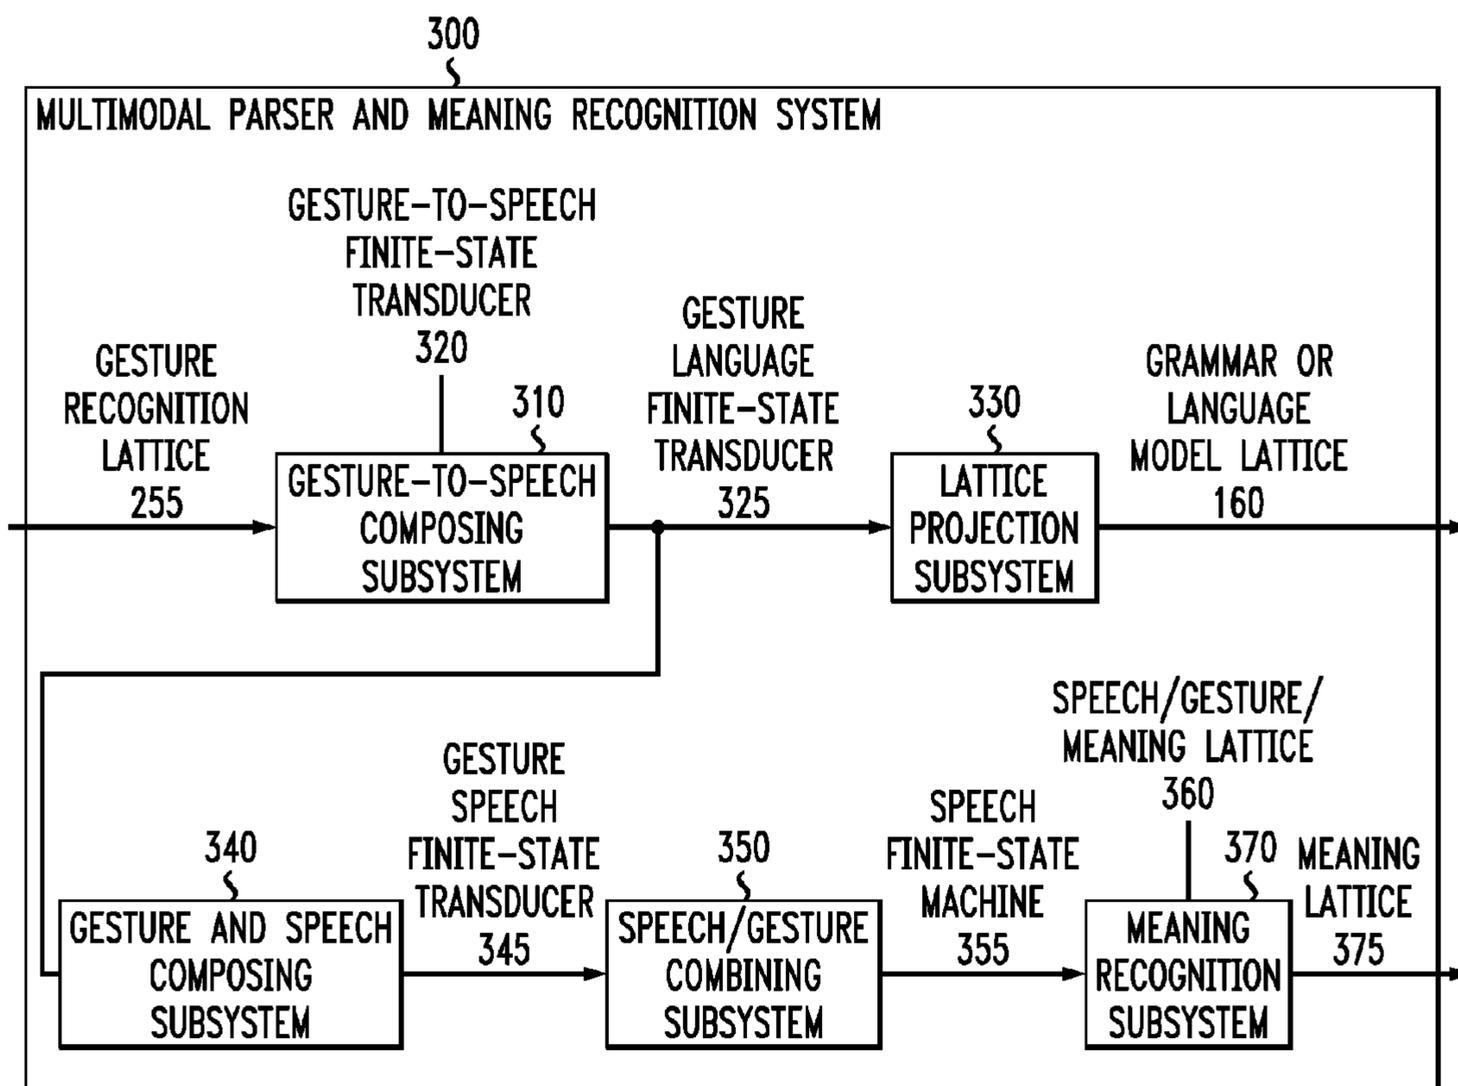


FIG. 6

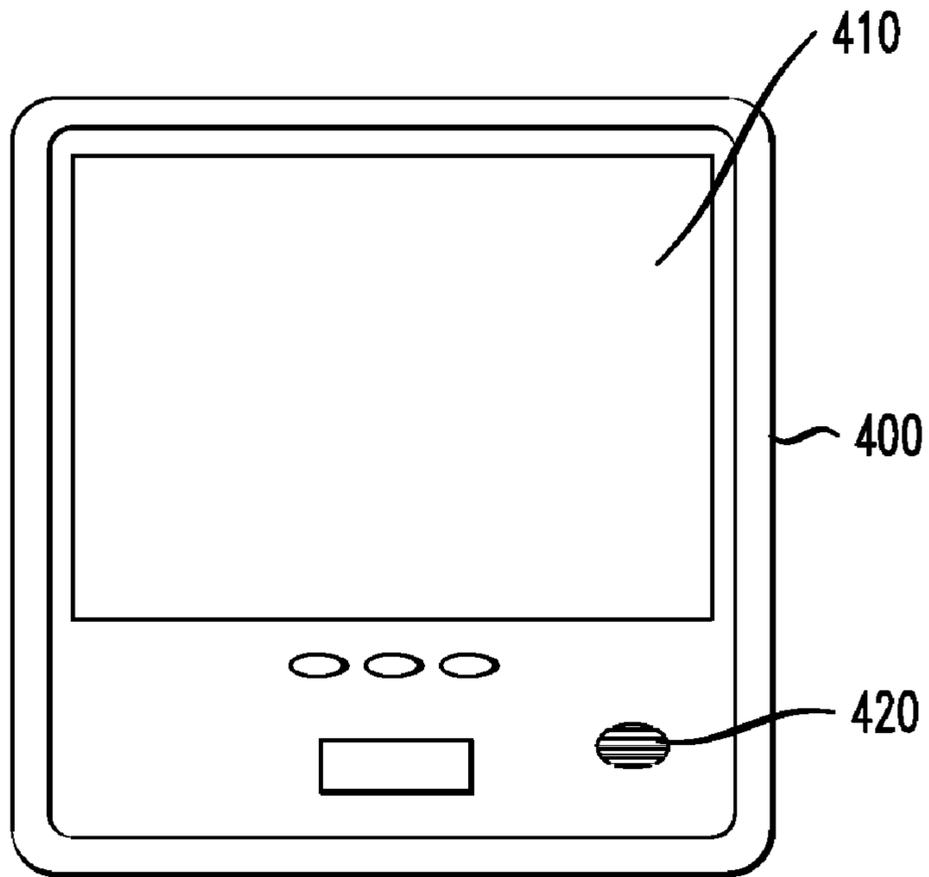


FIG. 8

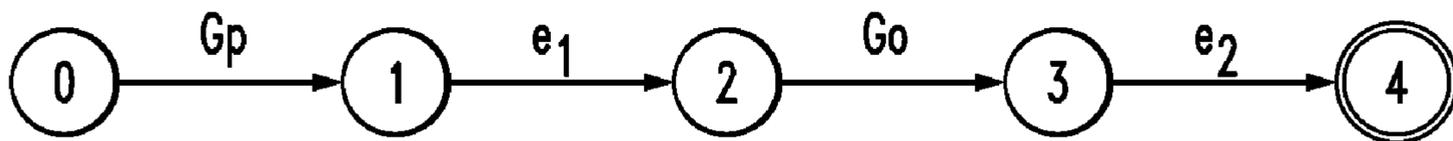
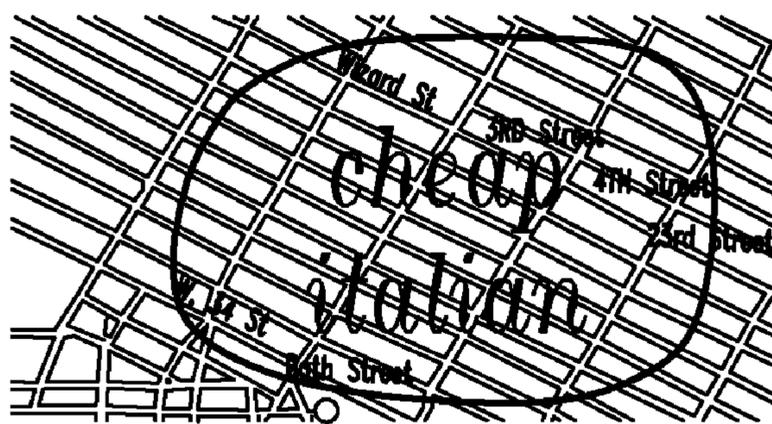


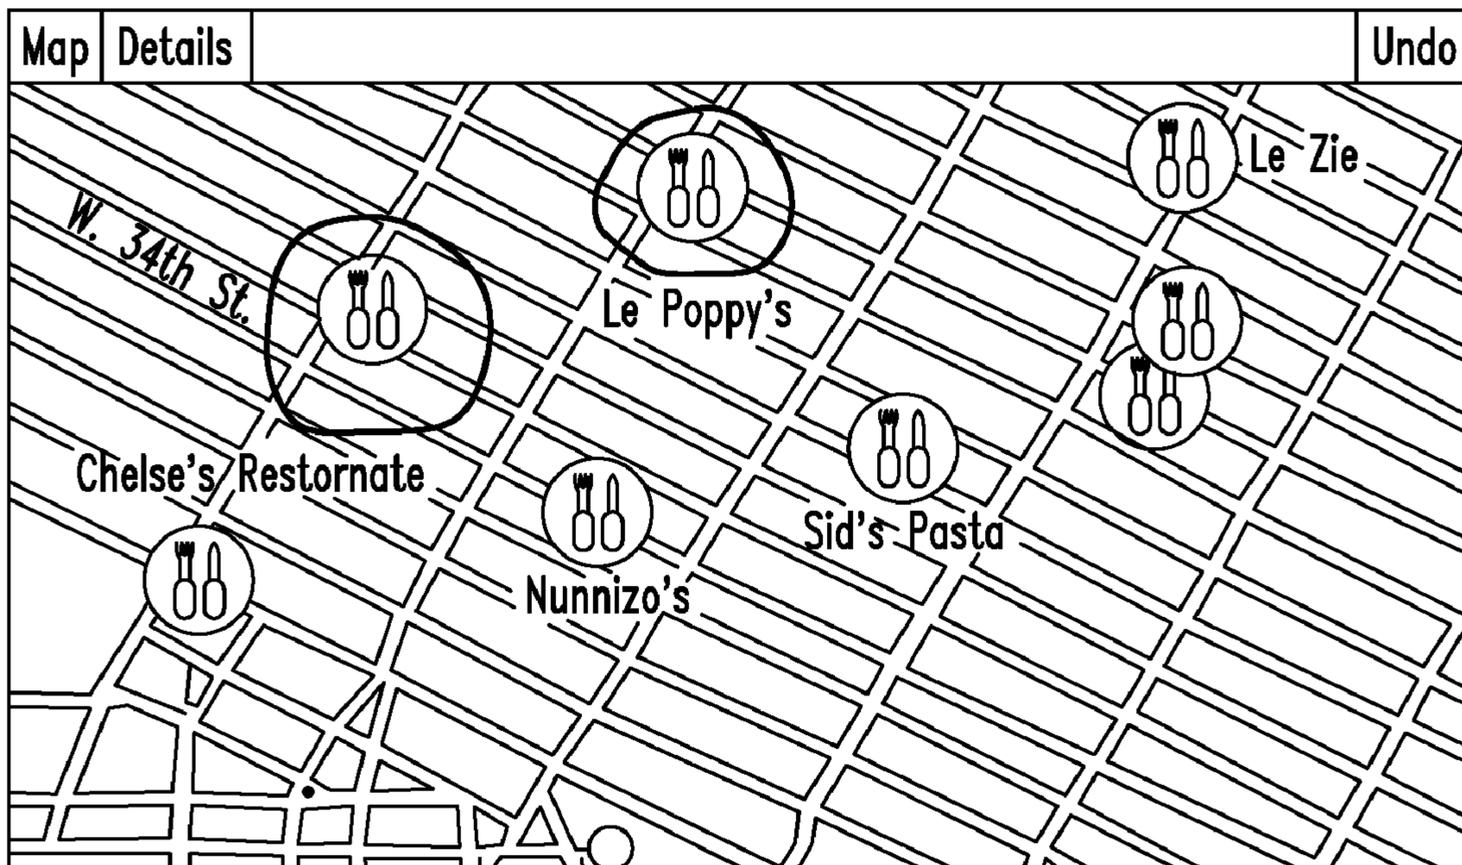
FIG. 9



*FIG. 7*

S	→	COMMAND
COMMAND	→	Show:eps:<show> NP eps:eps:/show>
COMMAND	→	tell:eps:<info> me:eps:eps about:eps:eps DEICTICNP eps:eps:</info>
NP	→	eps:eps:<restaurant> CUISMOD restaurants:eps:eps LOCMOD eps:eps:</restaurant>
DEICTICNP	→	DDETSG SELECTION eps:1:eps RESTSG eps:eps:<restaurant> Entry eps:eps</restaurant>
DEICTICNP	→	DDETPL SELECTION
	→	NUM RESTPL eps:eps:<restaurant>
	→	Entry eps:eps</restaurant>
SELECTION	→	eps:area:eps eps:selection:eps
CUISMOD	→	eps:eps:<cuisine> CUISINE eps:eps:</cuisine>
CUISINE	→	italian:eps:italian
CUISINE	→	chinese:eps:chinese
LOCMOD	→	eps:eps:<location> LOCATION eps:eps:</location>
LOCMOD	→	eps:eps:eps
LOCATION	→	in:eps:eps this:G:eps area:area:eps eps:location:eps Entry
LOCATION	→	along:eps:eps this:G:eps route:line:eps eps:location:eps Entry
DDETSG	→	this:G:eps
DDETPL	→	these:G:eps
NUM	→	two:2:eps
NUM	→	three:3:eps
RESTSG	→	restaurant:restaurant:eps
RESTPL	→	restaurants:restaurant:eps
Entry	→	eps:SEM:SEM

FIG. 10





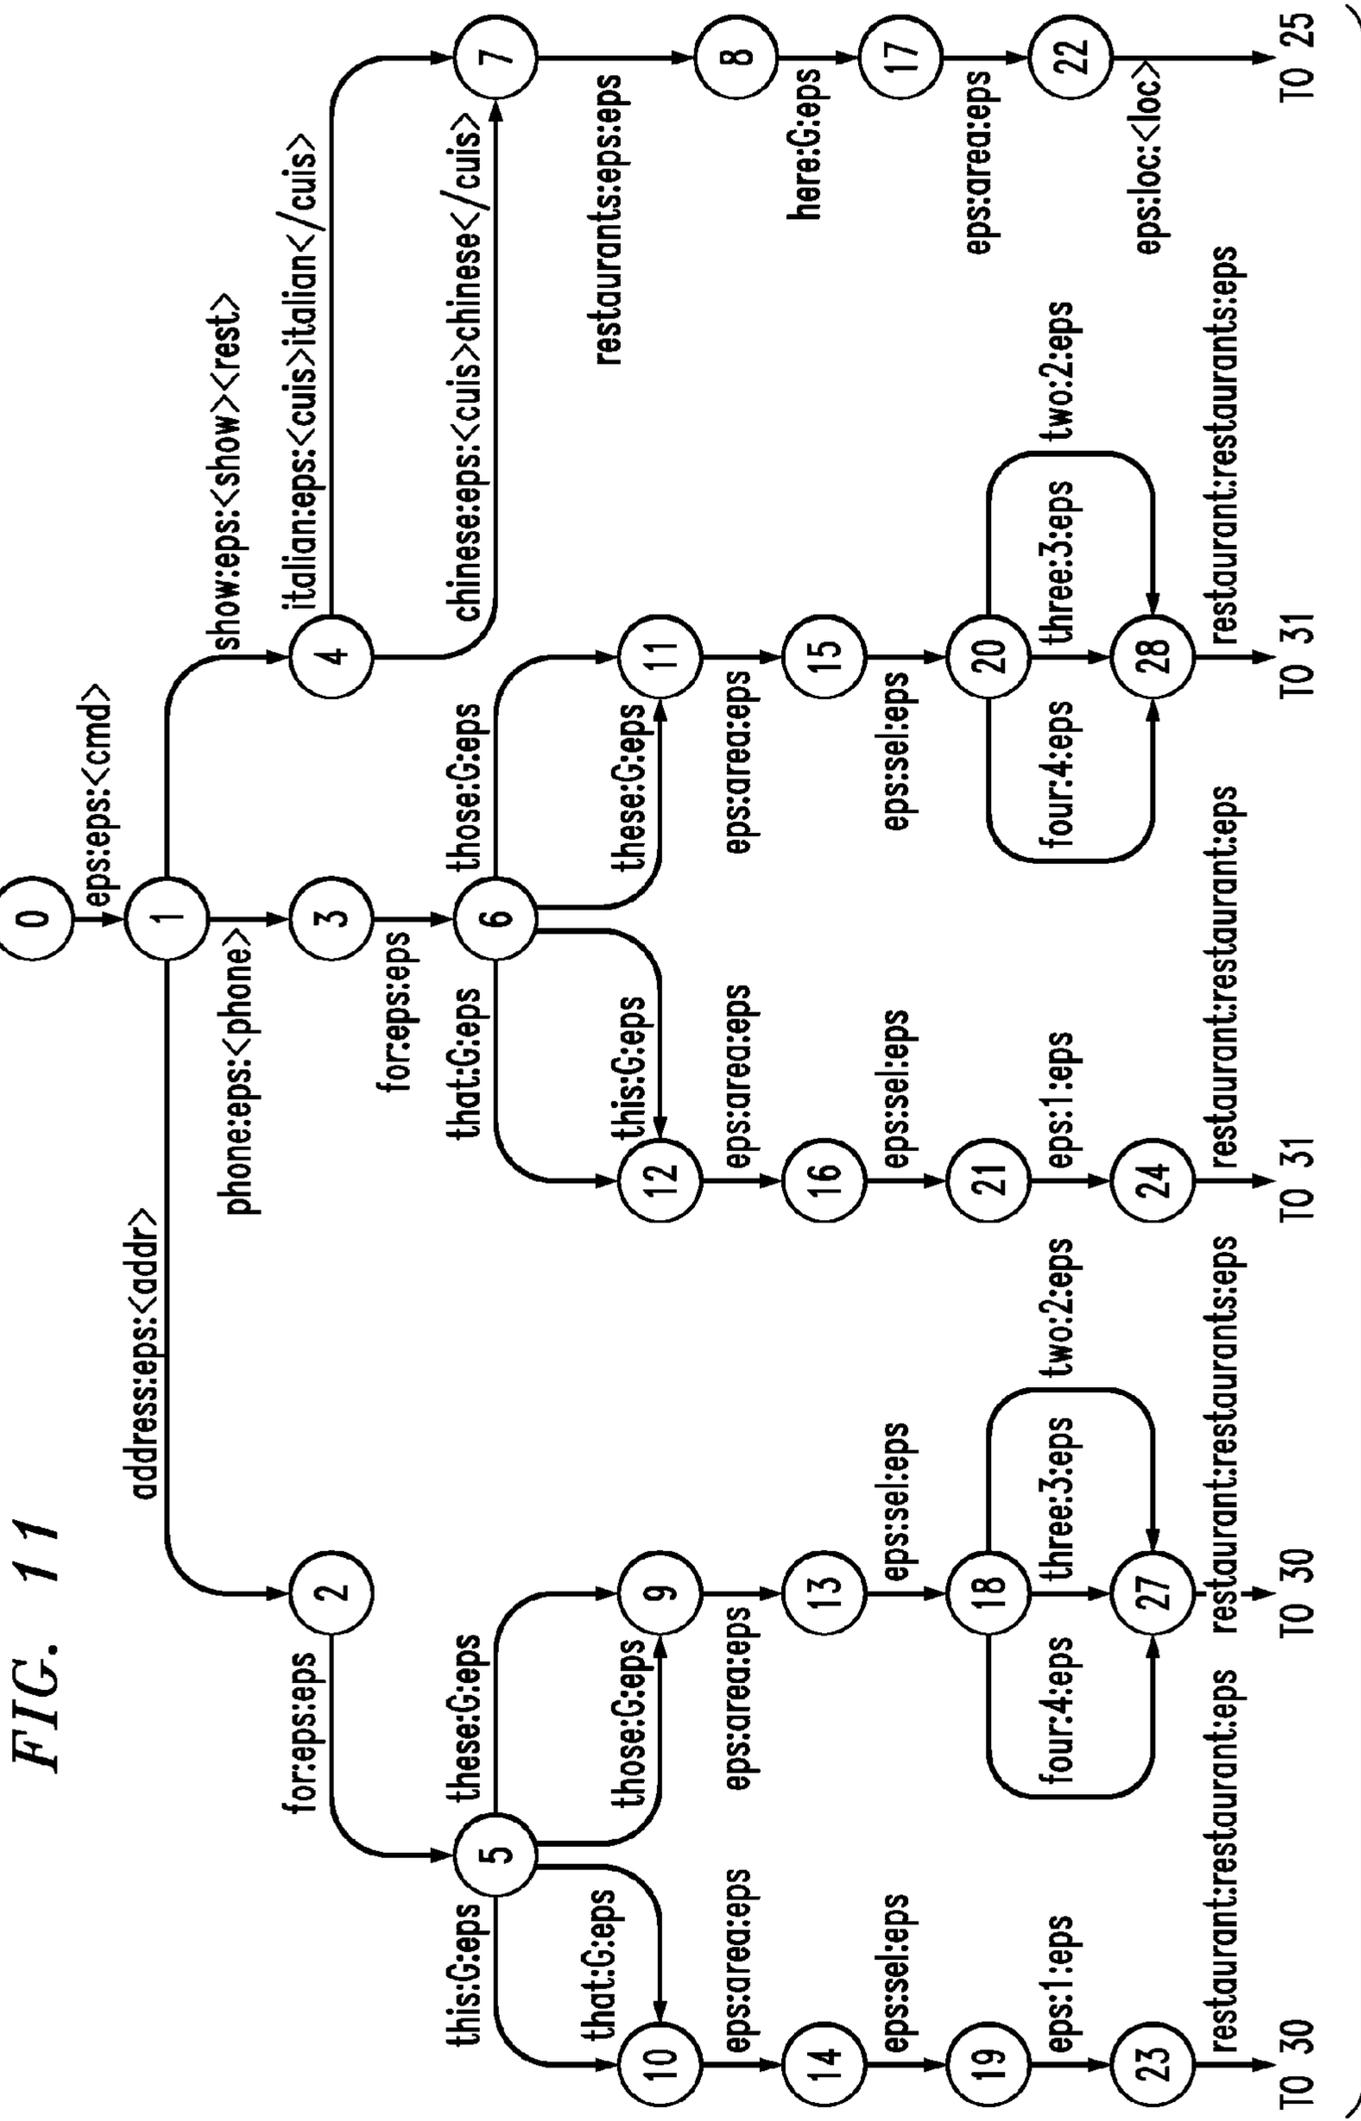
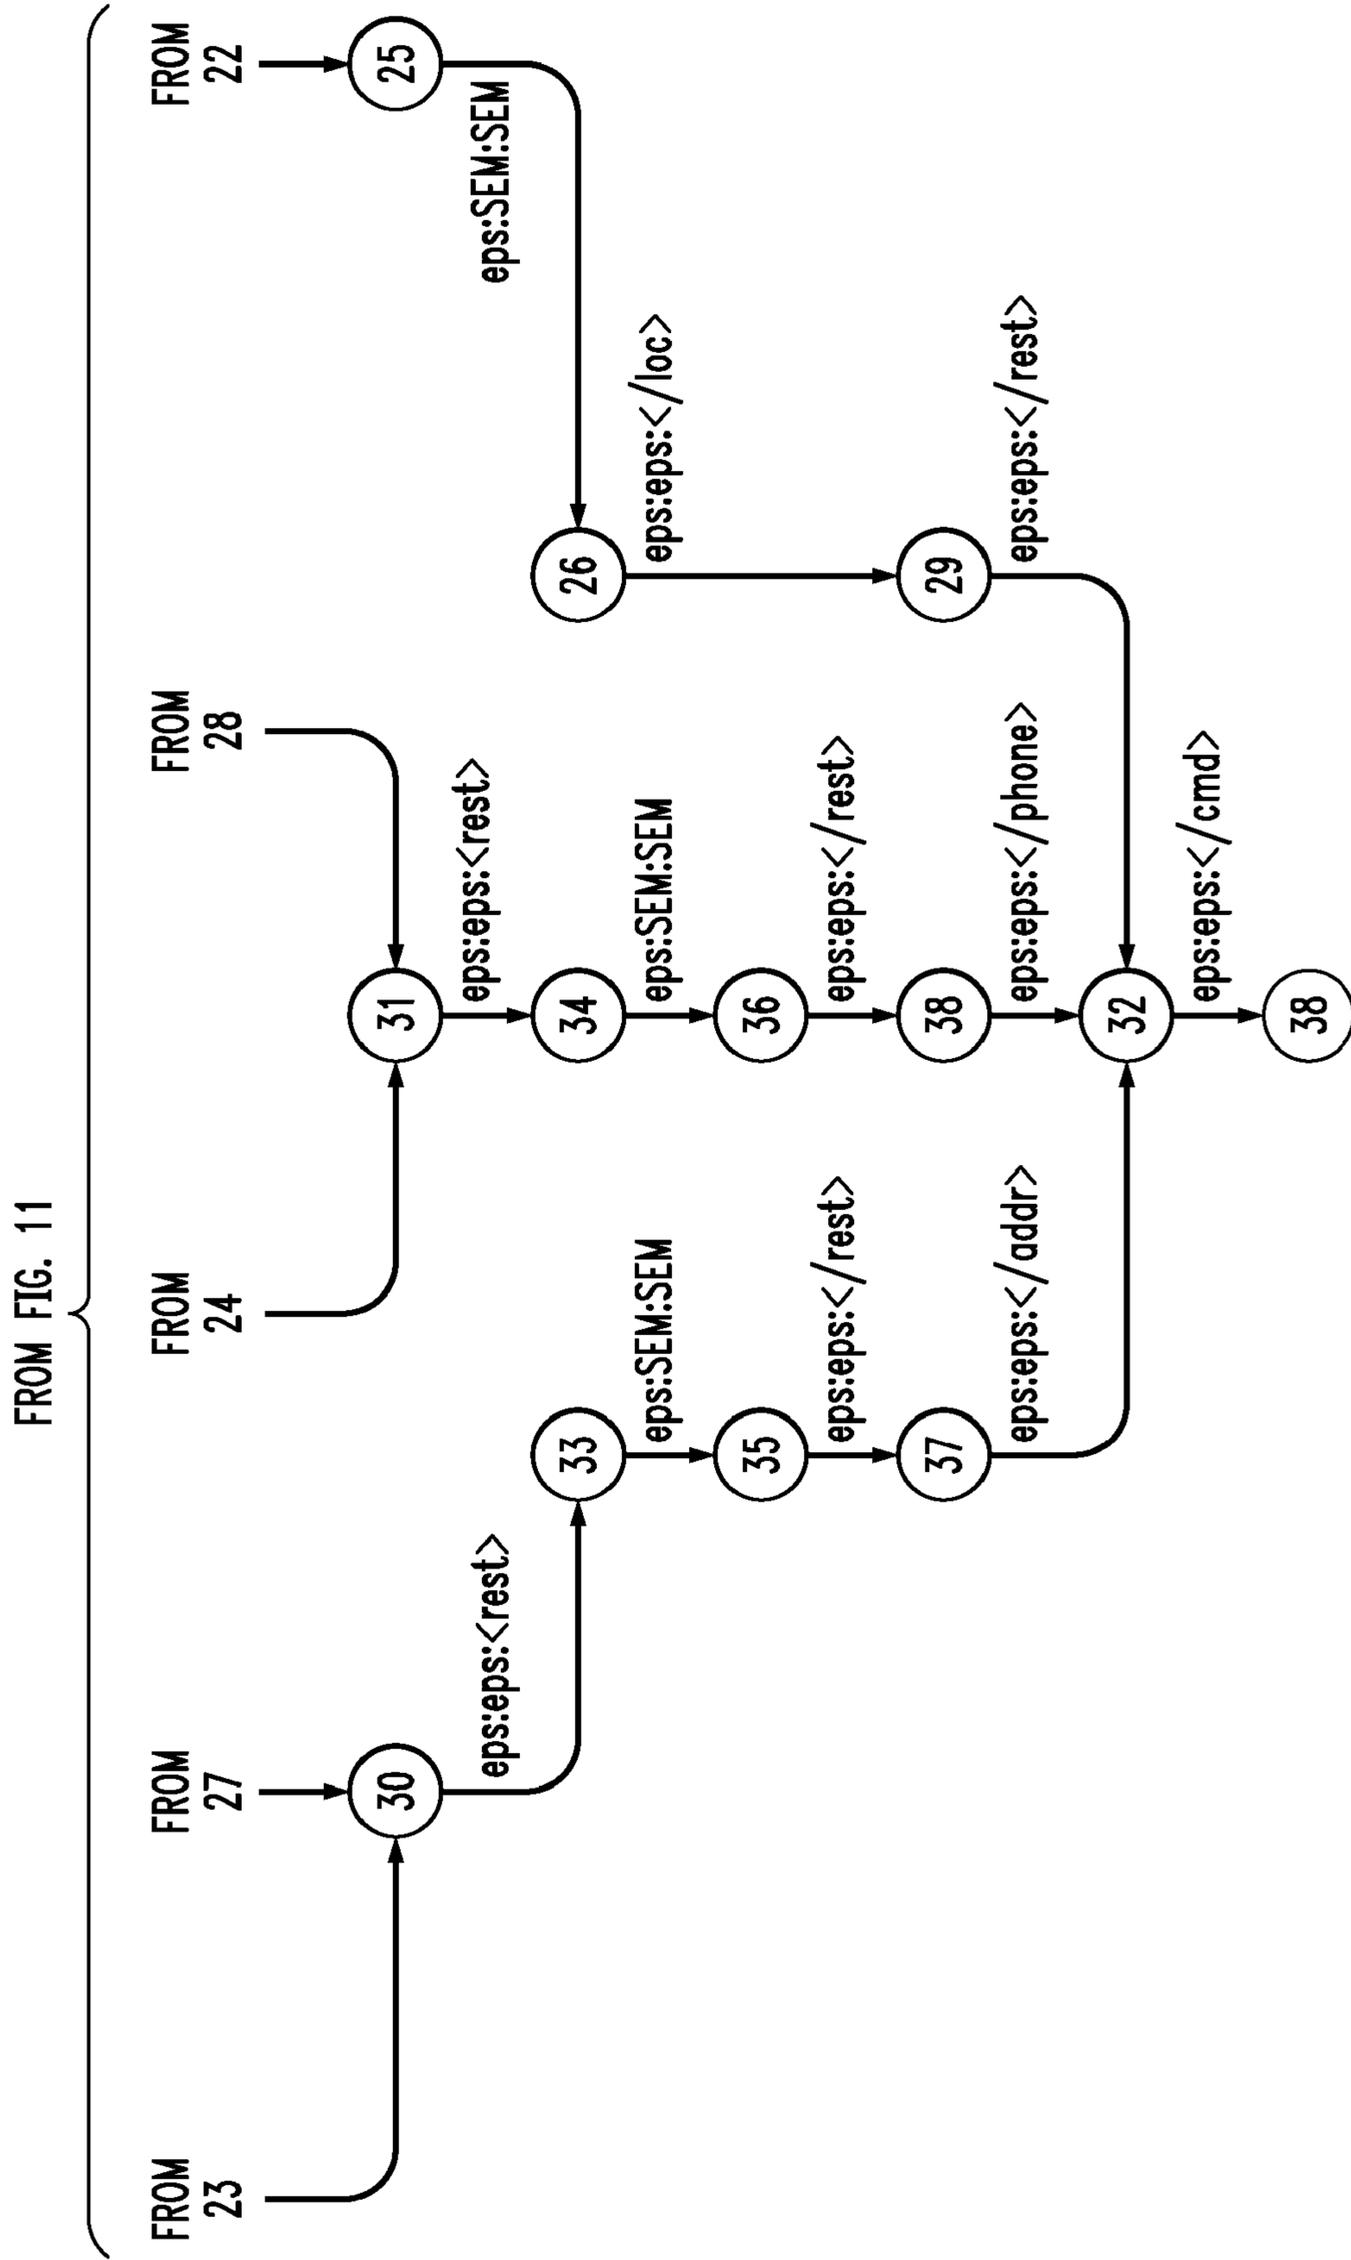


FIG. 11

TO FIG. 11 cont.

FIG. 11 cont.



*FIG. 12*

S eps:eps:<cmd> \_CMD eps:eps:</cmd>

\_CMD \_PHONECMD  
\_CMD \_ADDRESSCMD  
\_CMD \_SHOWRESTAURANTCMD

\_PHONECMD phone:eps:<phone> for:eps:eps \_DEICTICNP eps:eps:</phone>  
\_ADDRESSCMD address:eps:<addr> for eps:eps DEICTICNP eps:eps:</addr>

\_SHOWRESTAURANTCMD show:eps:eps:<show><rest> \_CUISINE restaurants:eps:eps  
\_AREALLOCATION eps:eps:</rest>

\_DEICTICNP \_DDETSG eps:area:eps: eps:sel:eps: eps:1:eps \_RESTSG  
eps:eps:<rest> eps:SEM:SEM eps:eps:</rest>

\_DEICTICNP \_DDETPL eps:area:eps eps:sel:eps \_NUM \_RESTPL  
eps:eps:<rest> eps:SEM:SEM eps:eps:</rest>

\_DEICTICNP \_DDETSG eps:area:eps eps:sel:eps eps:1:eps \_RESTSG  
eps:eps:<rest> eps:SEM:SEM eps:eps:</rest> and:eps:eps \_DEICTICNP

DEICTICNP \_DDETPL eps:area:eps eps:sel:eps \_NUM \_RESTPL  
eps:eps:<rest> eps:SEM:SEM eps:eps:</rest> and:eps:eps \_DEICTICNP

\_DDETPL these:G:eps  
\_DDETPL those:G:eps  
\_DDETPL this:G:eps  
\_DDETPL that:G:eps

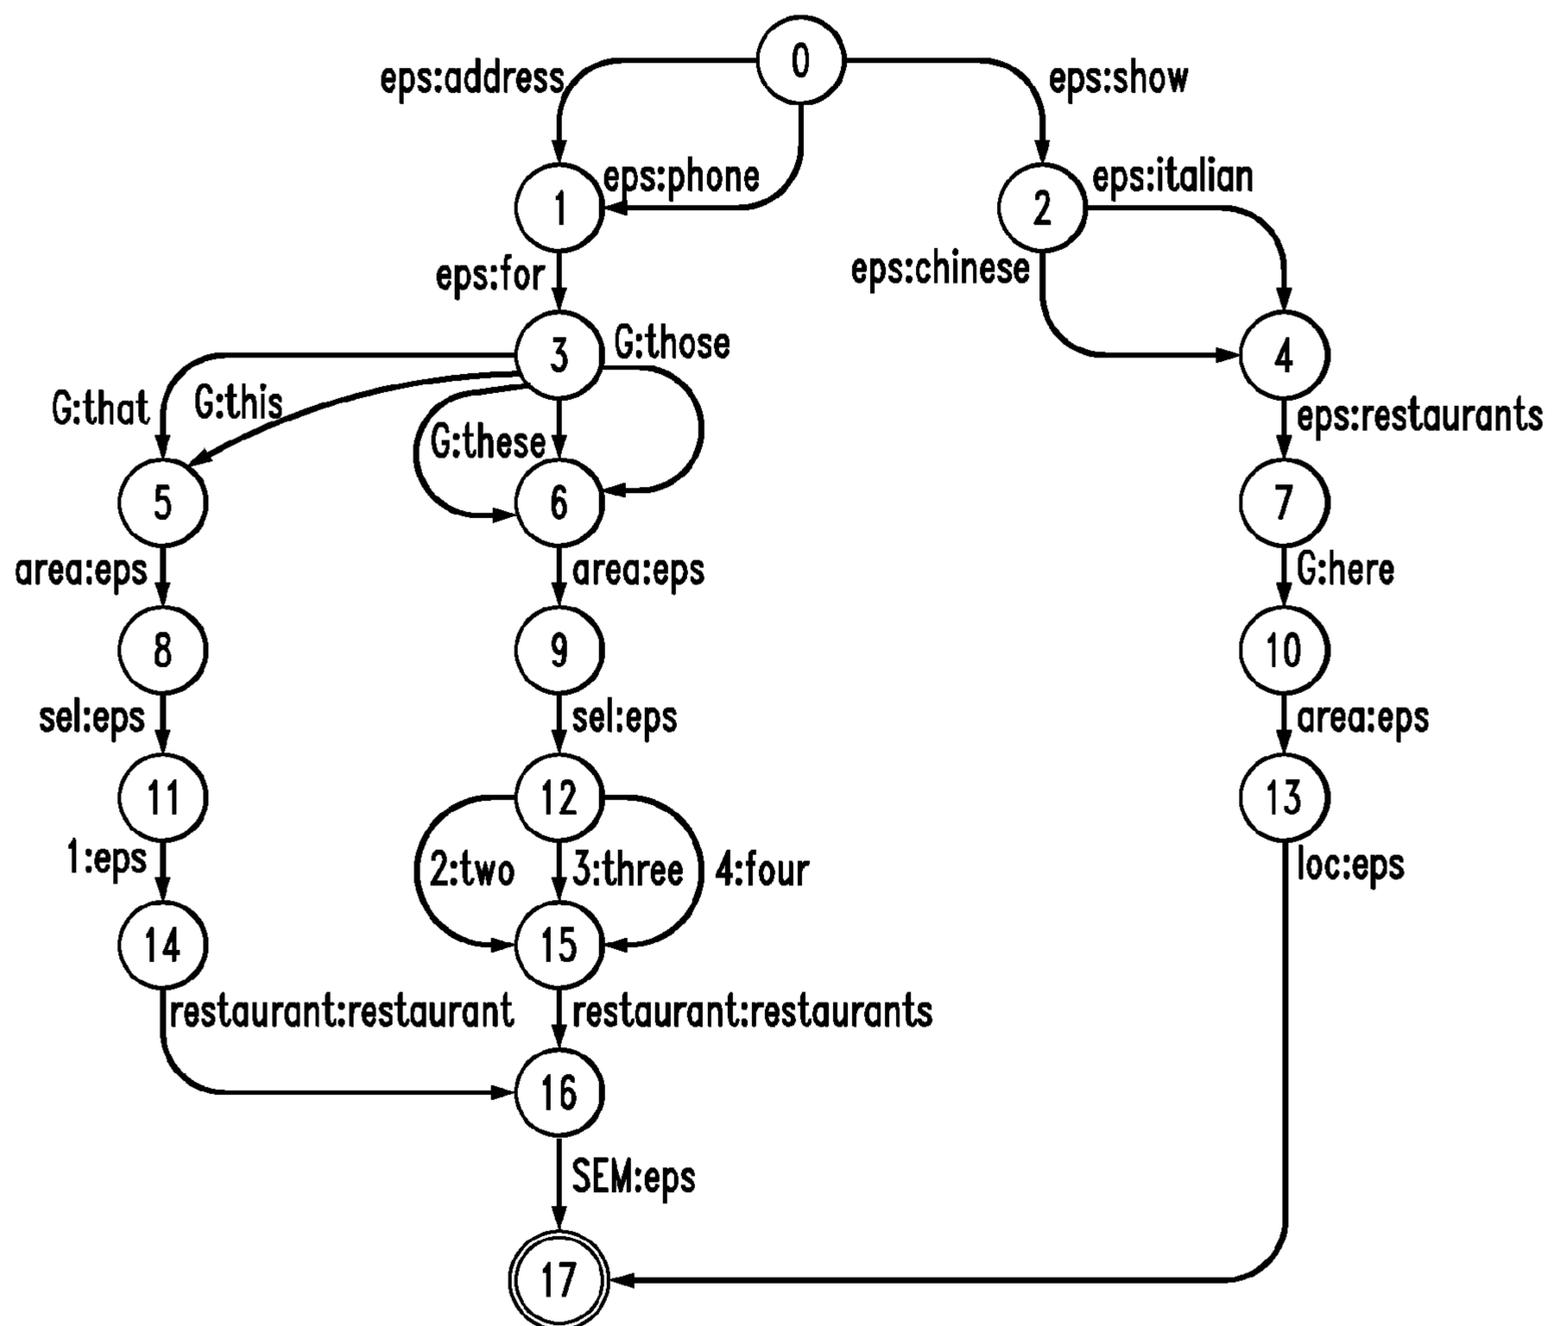
\_AREALLOCATION here:G:eps eps:area:eps eps:loc:<loc>  
eps:SEM:SEM eps:eps:</loc>

\_RESTPL restaurants:restaurant:eps  
\_RESTSG restaurant:restaurant:eps

\_NUM three:3:eps  
\_NUM two:2:eps  
\_NUM four:4:eps

\_CUISINE chinese:eps:<cuis>chinese</cuis>  
\_CUISINE italian:eps:<cuis>italian</cuis>

FIG. 13



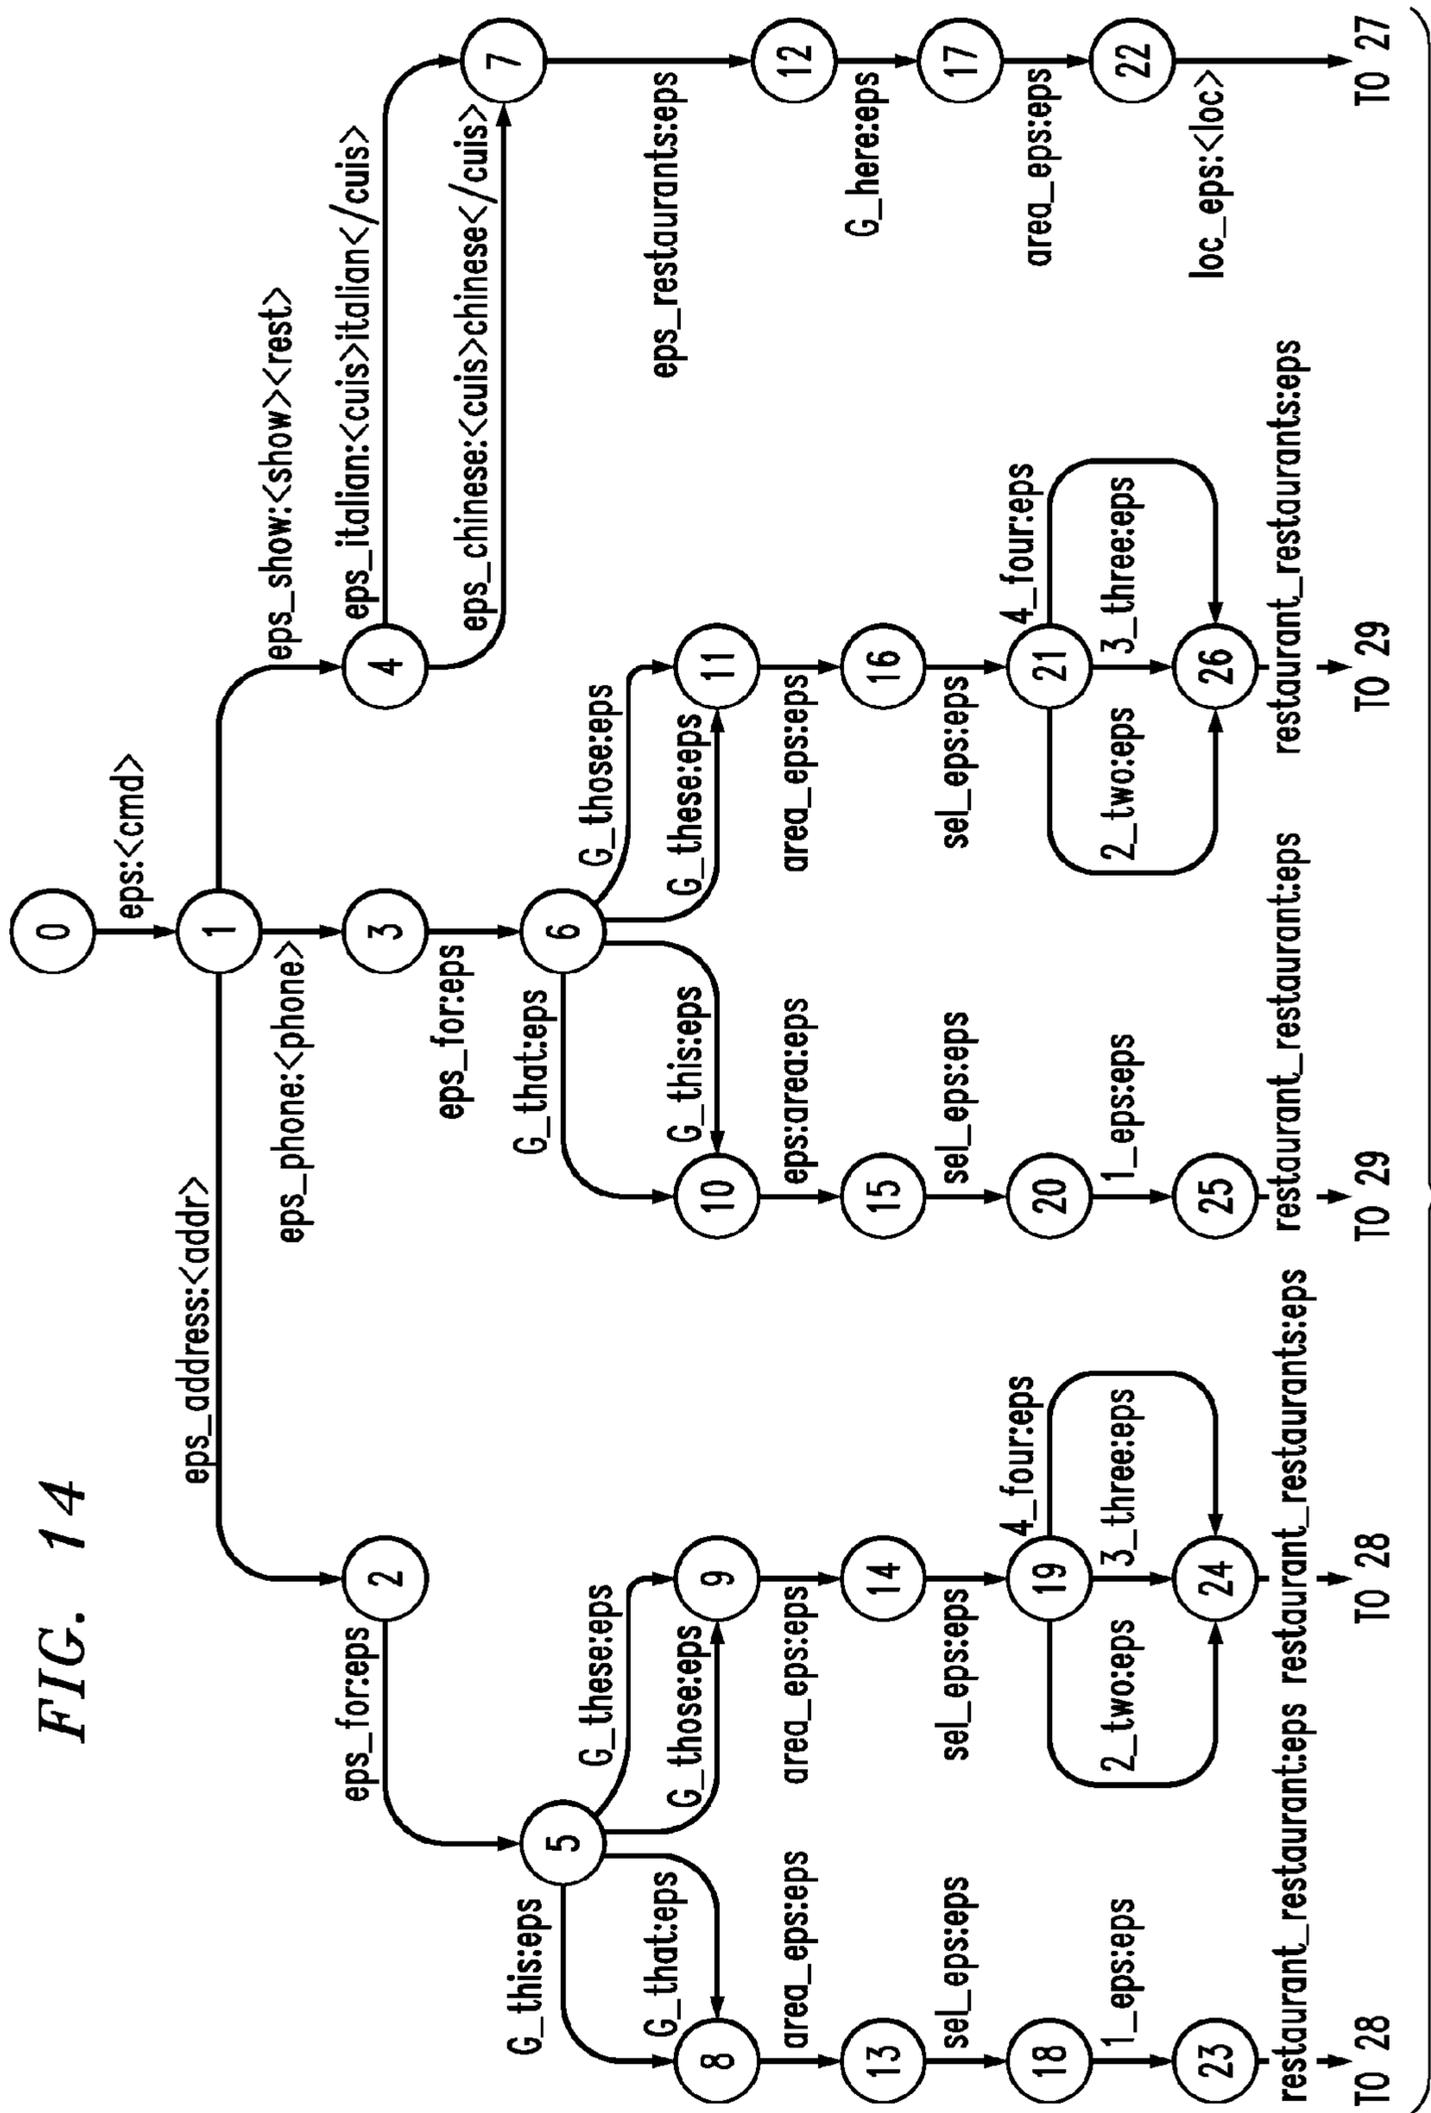


FIG. 14

TO FIG. 14 cont.

FIG. 14 cont.

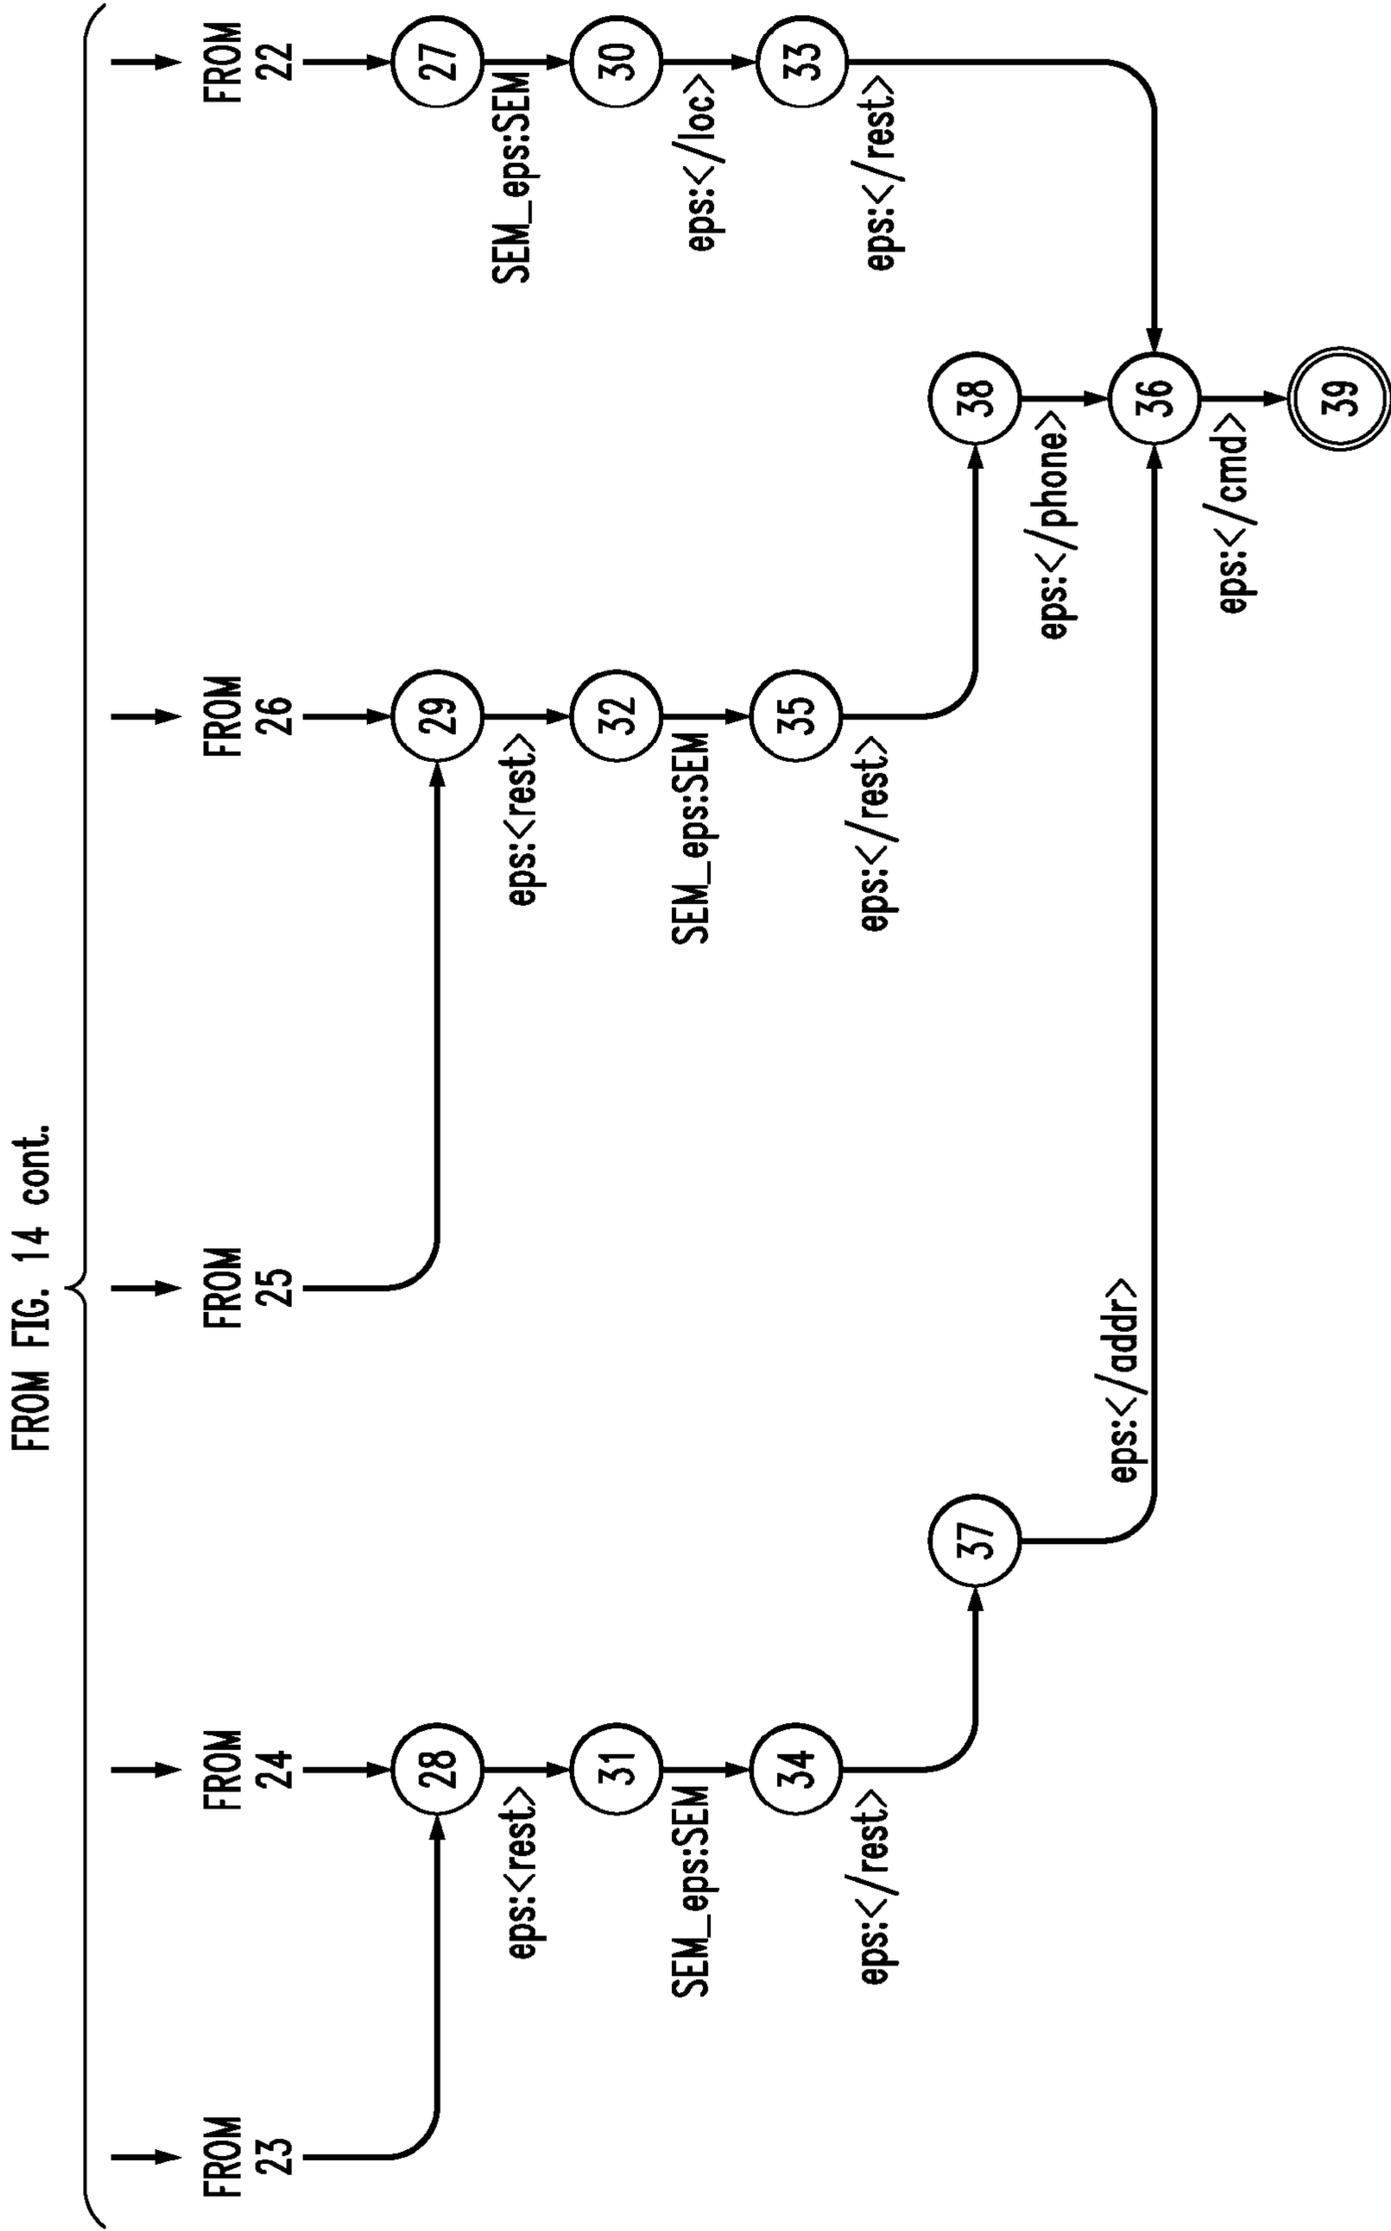


FIG. 15

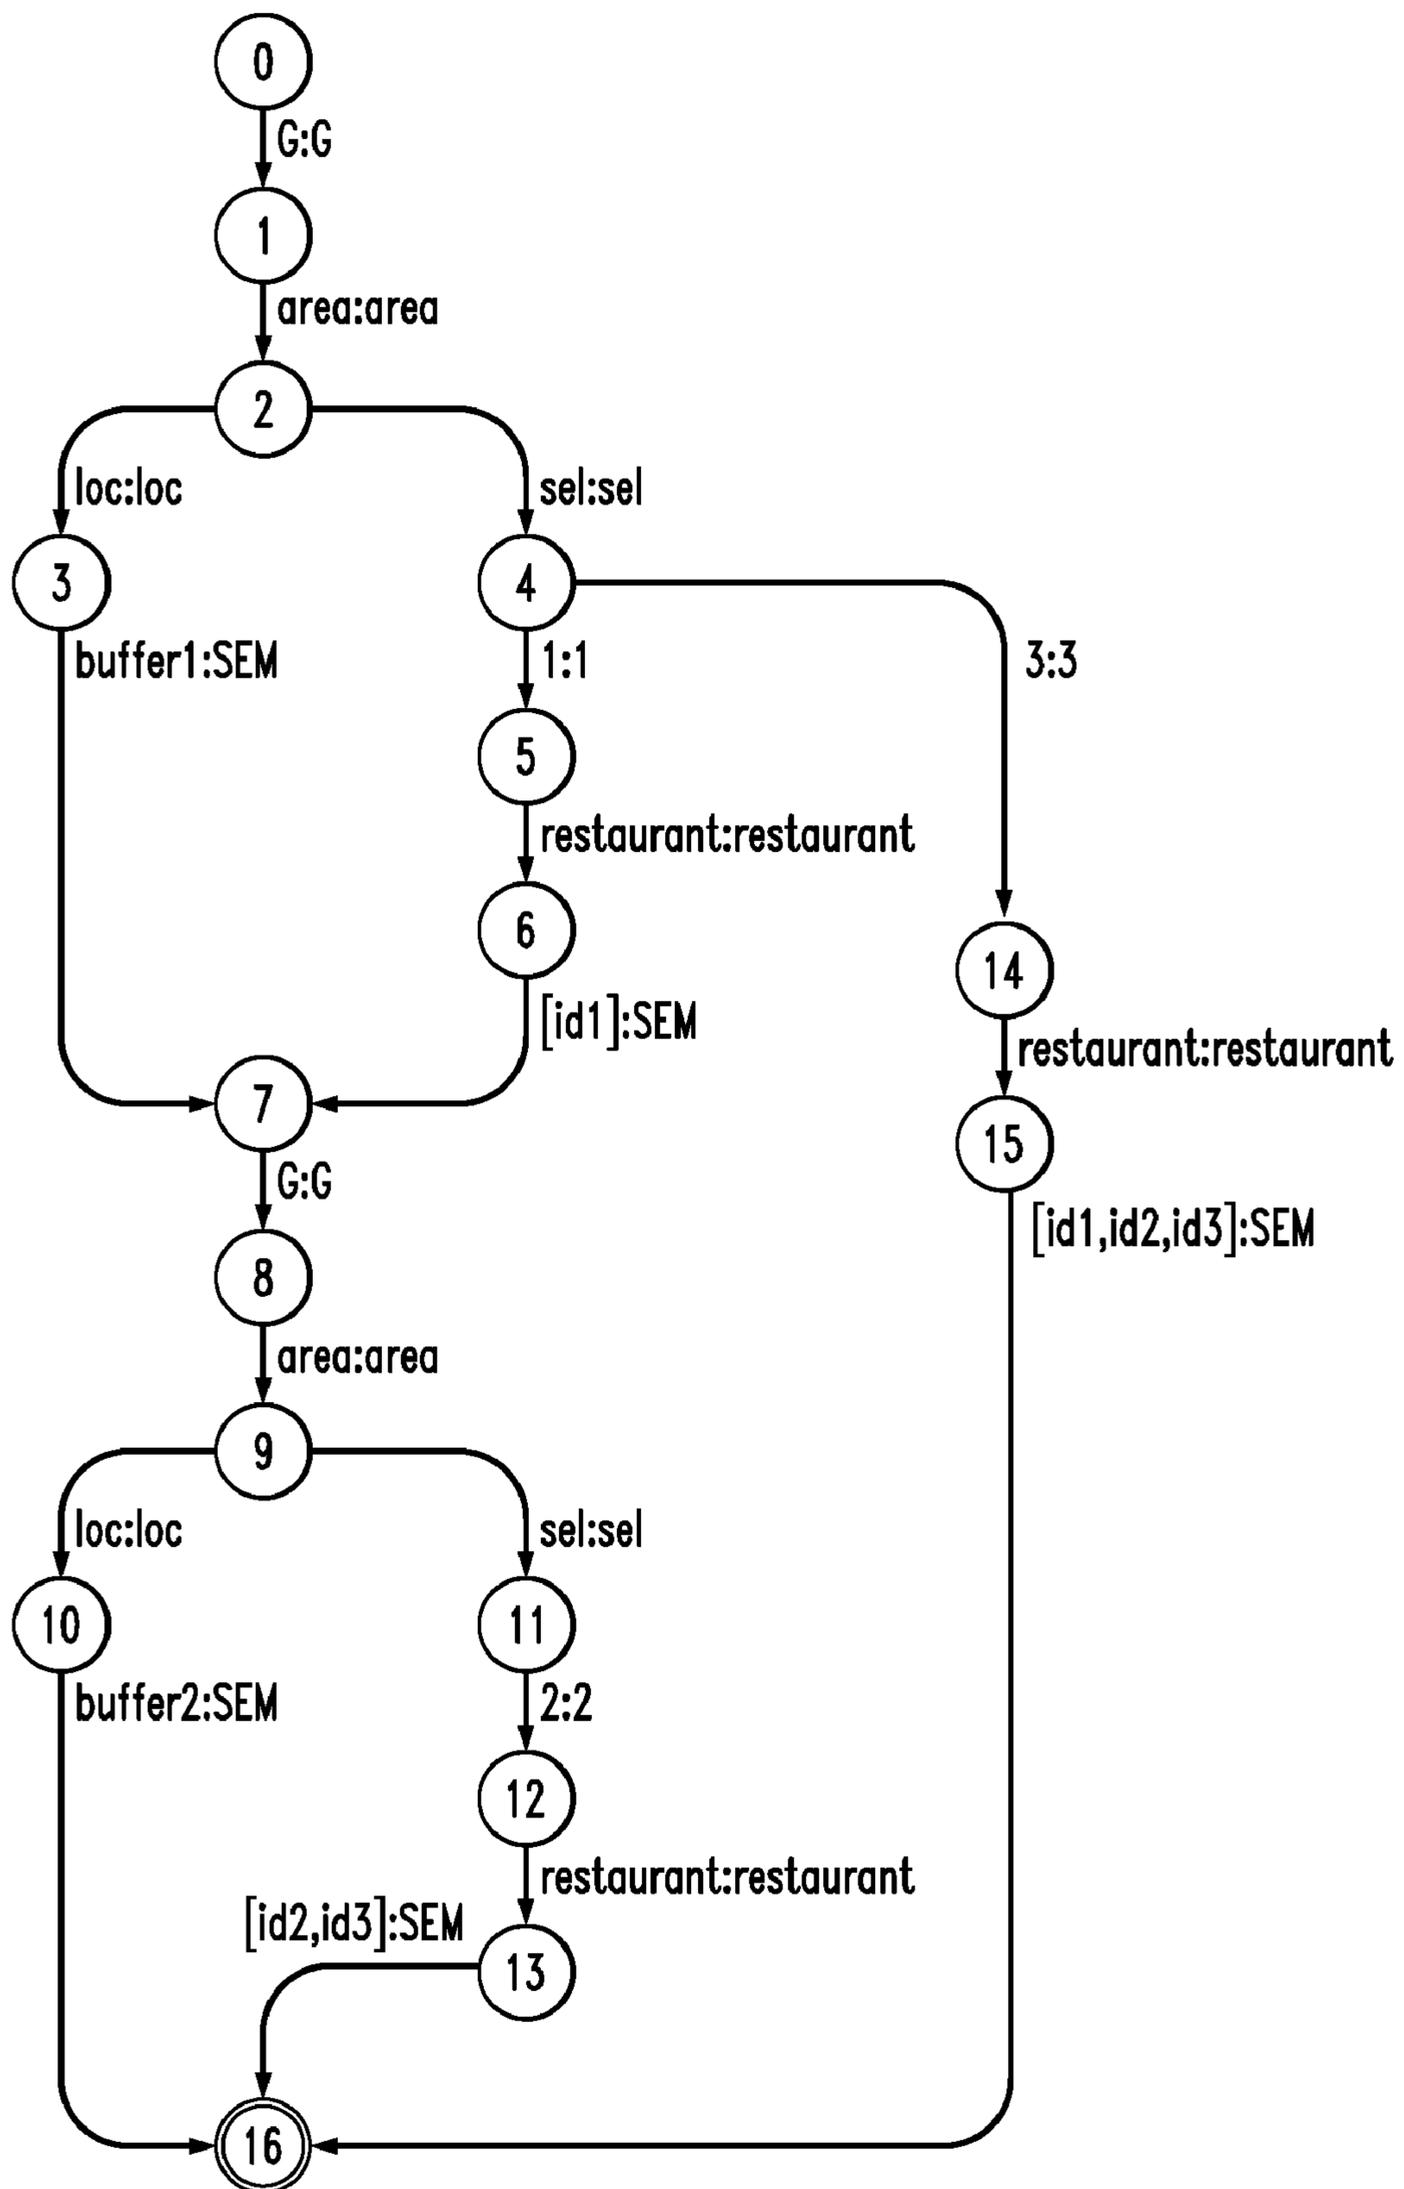


FIG. 16

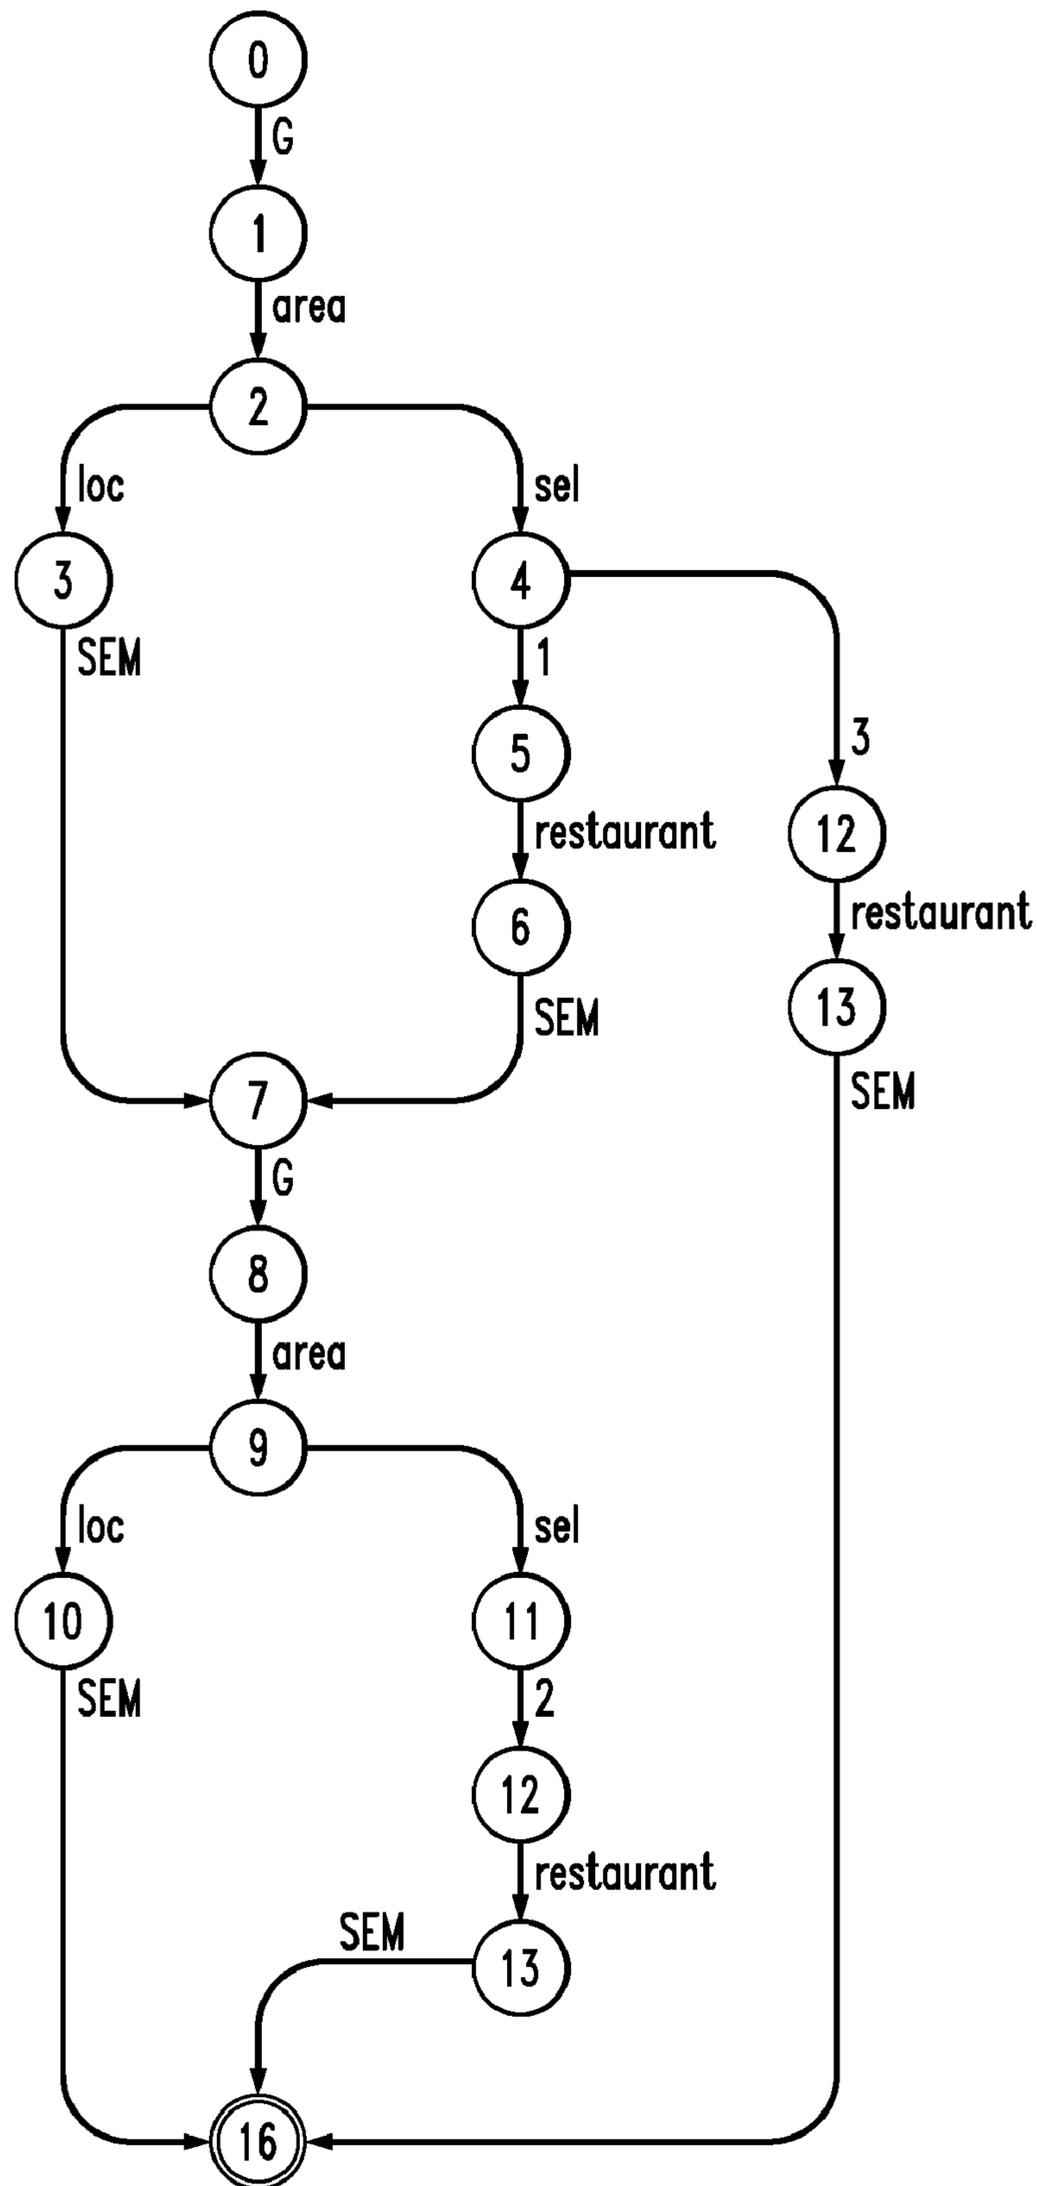




FIG. 17

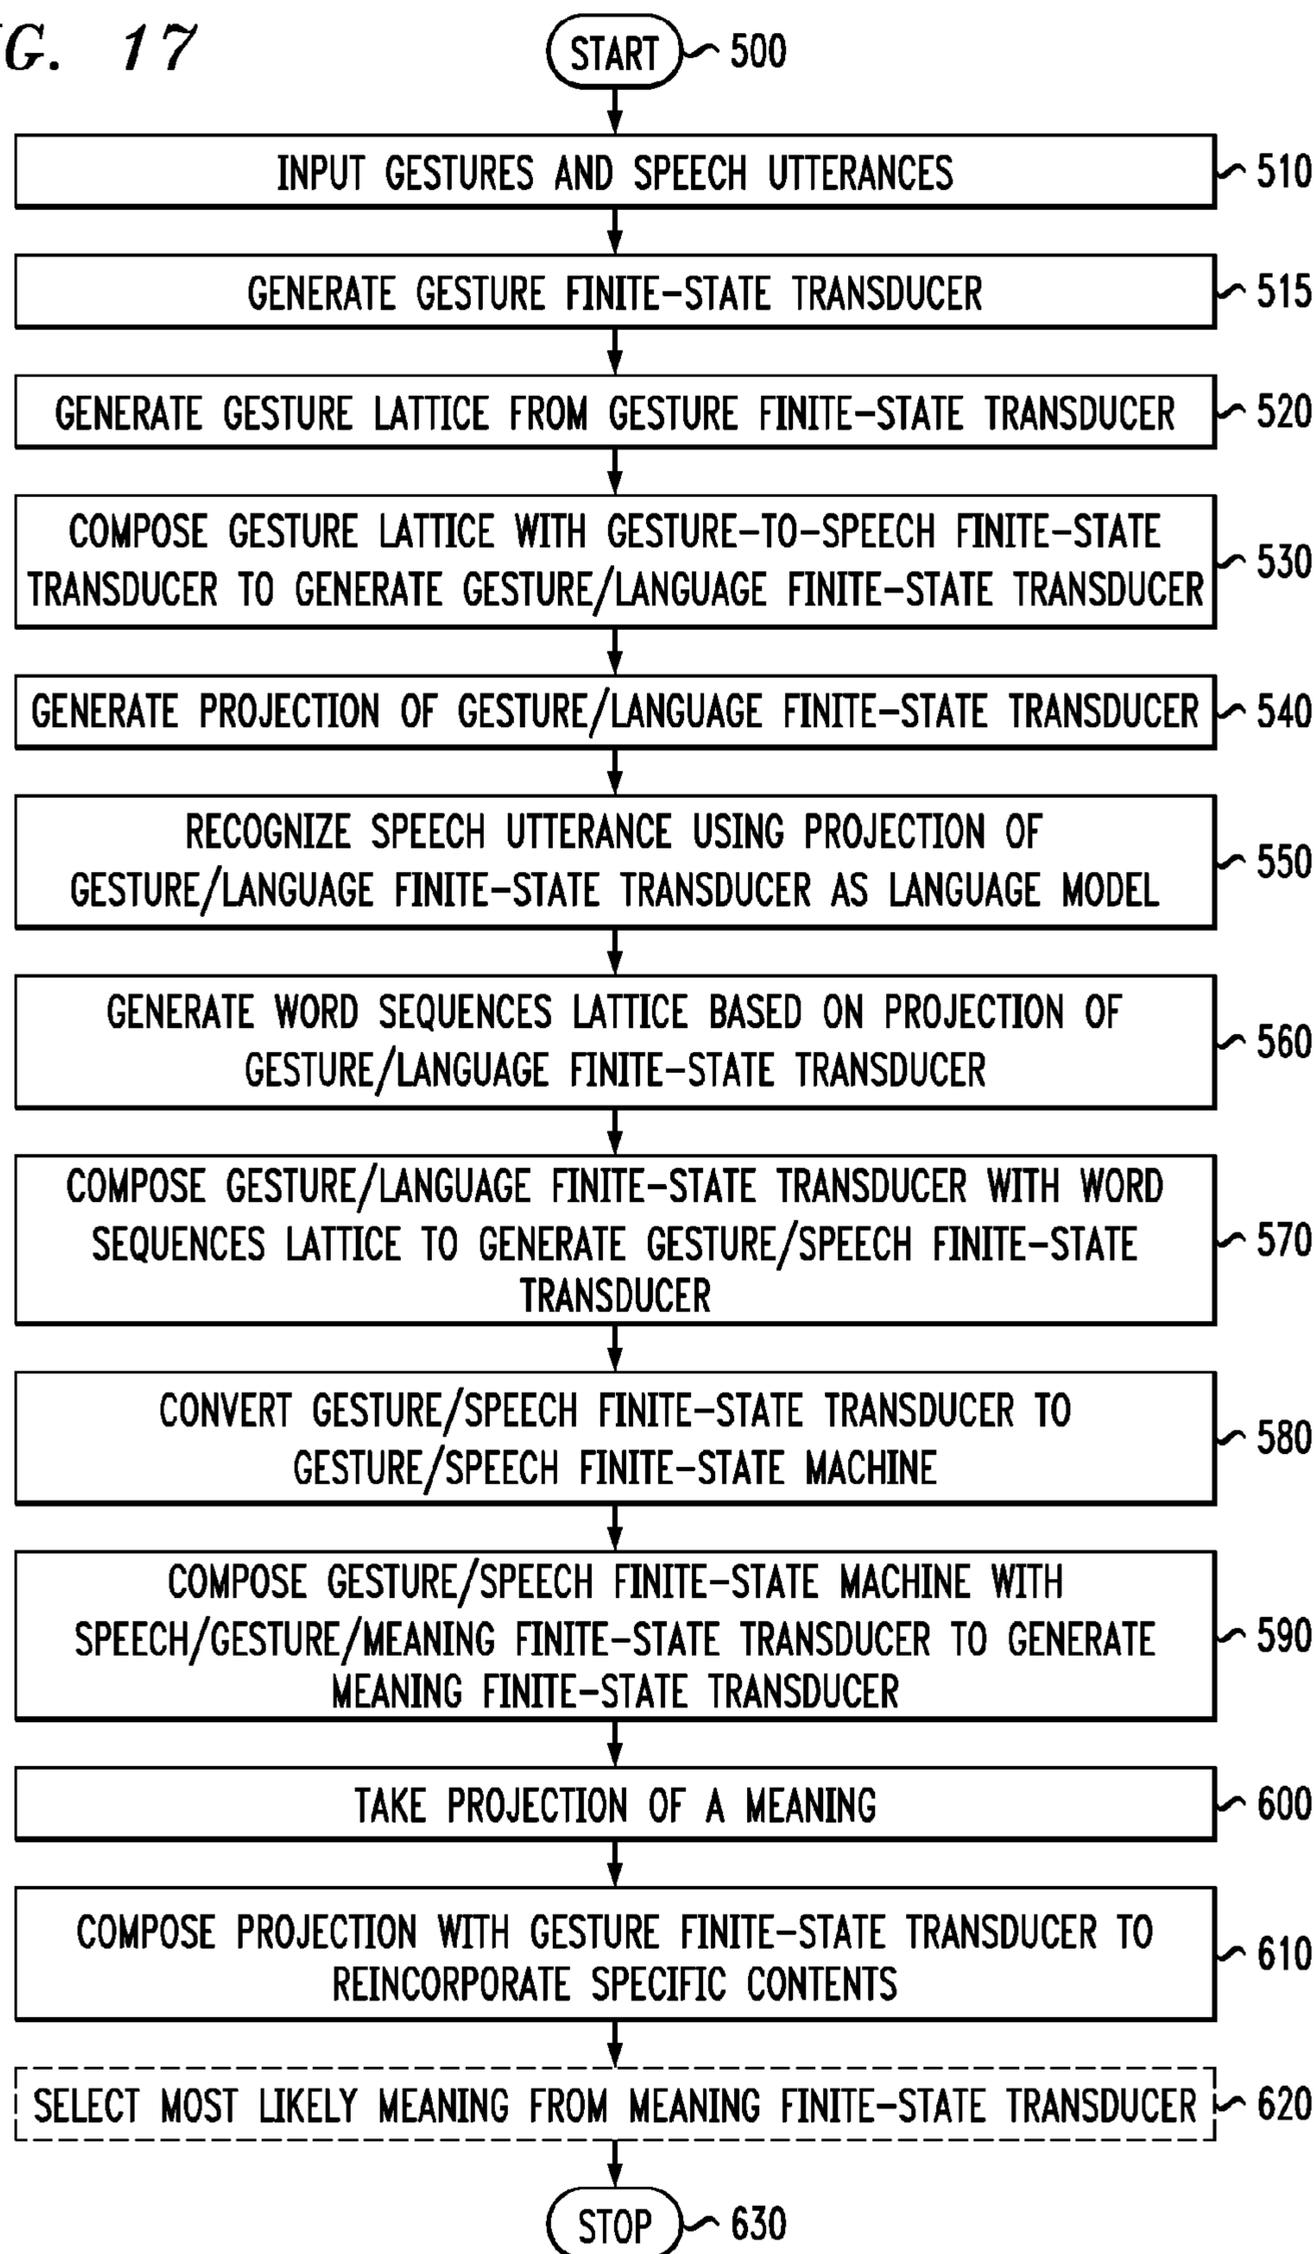


FIG. 18

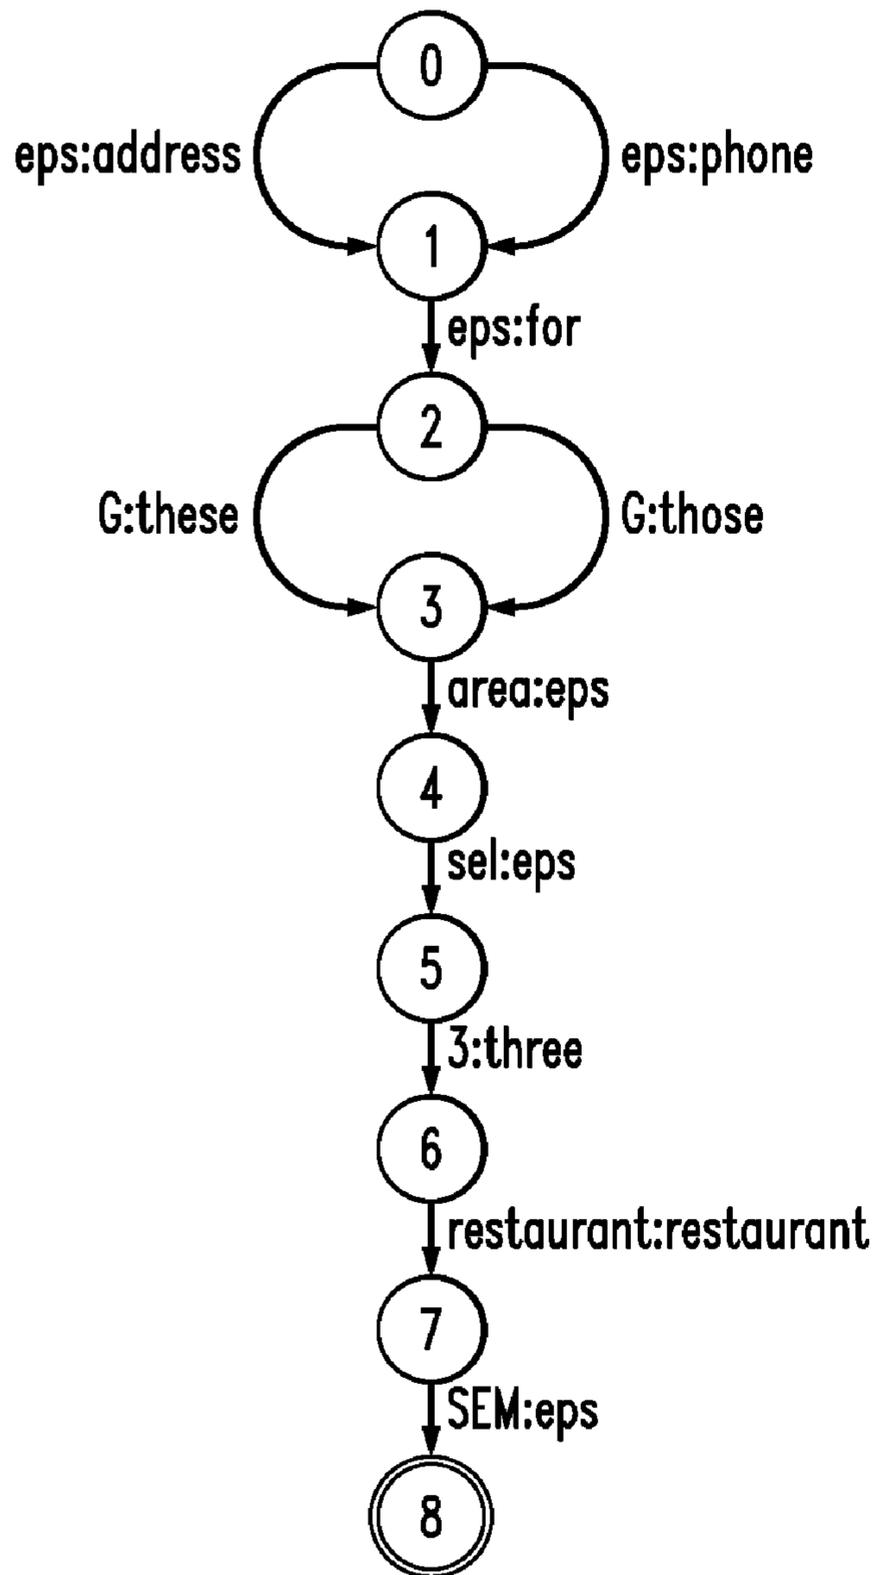


FIG. 19

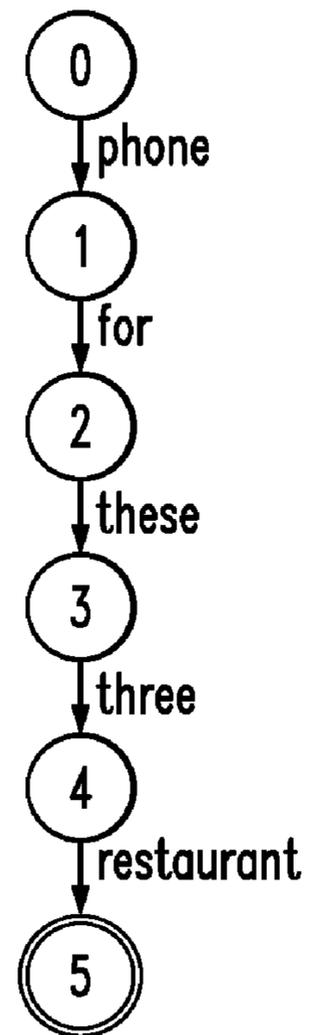


FIG. 20

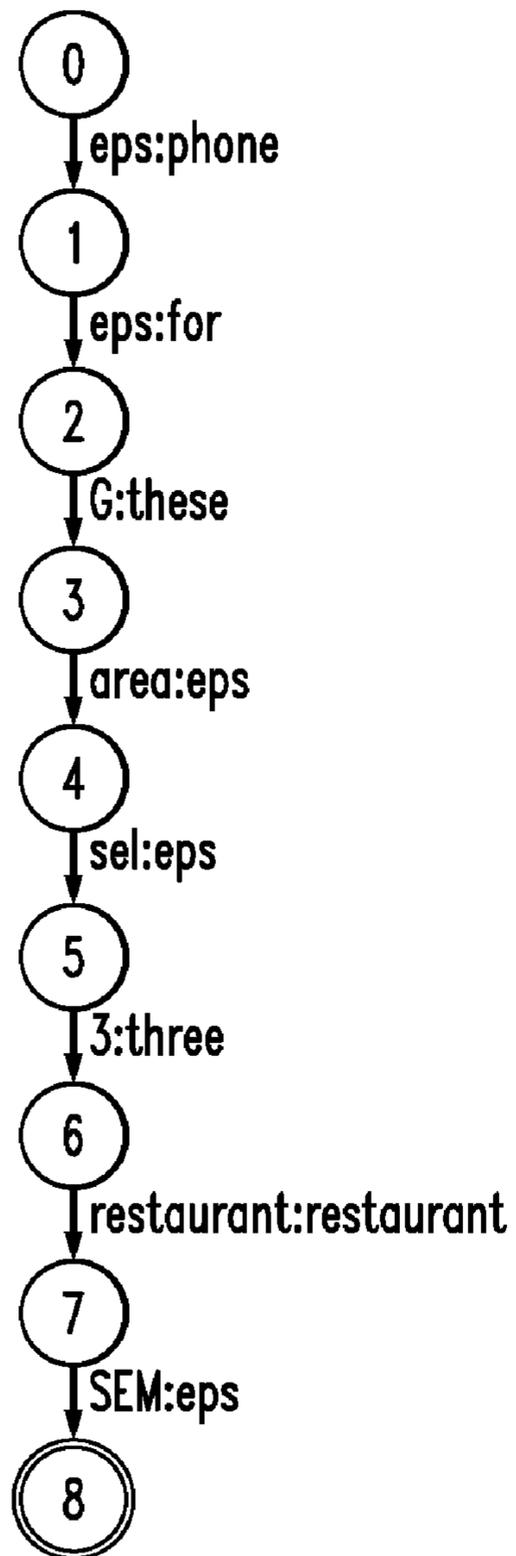


FIG. 21

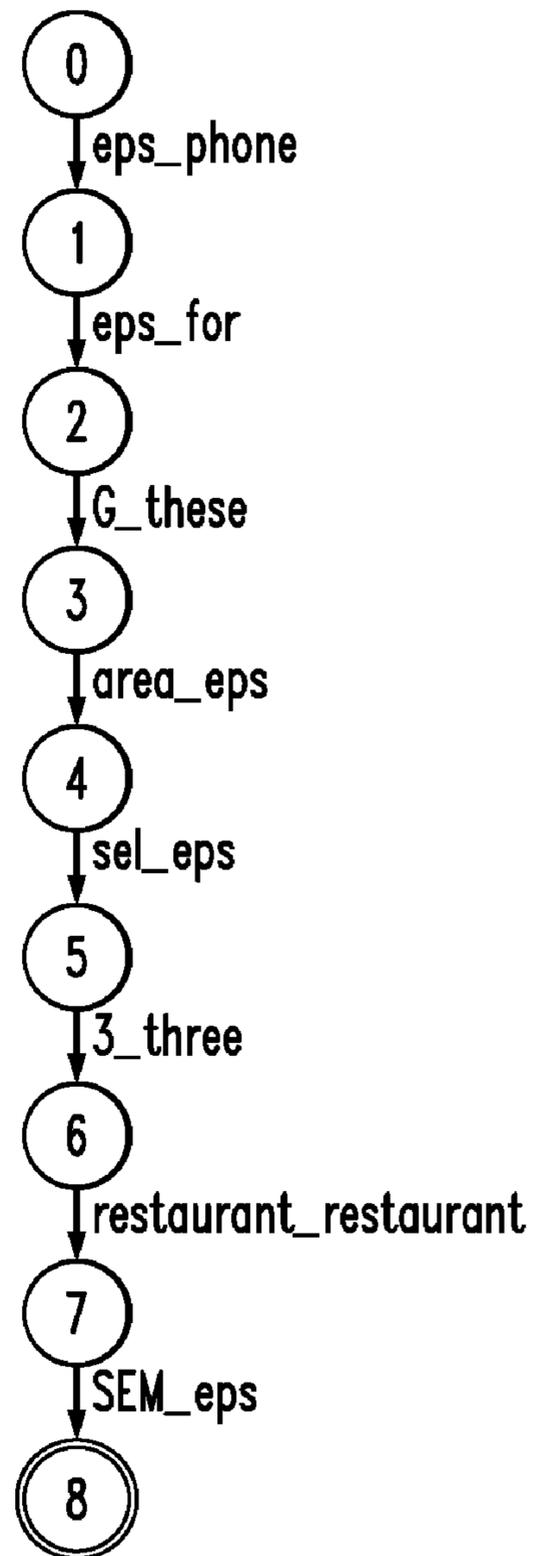
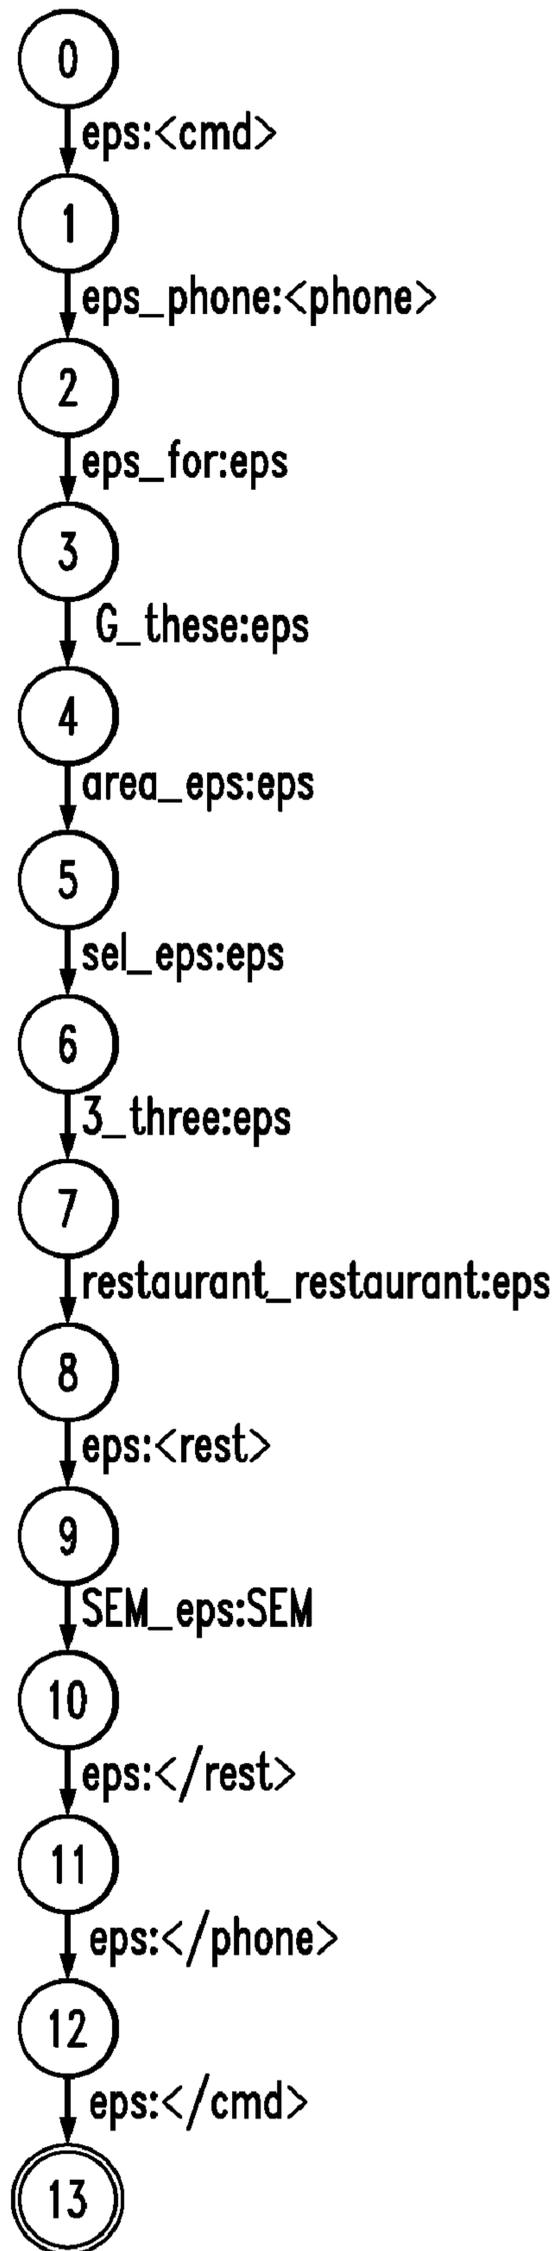
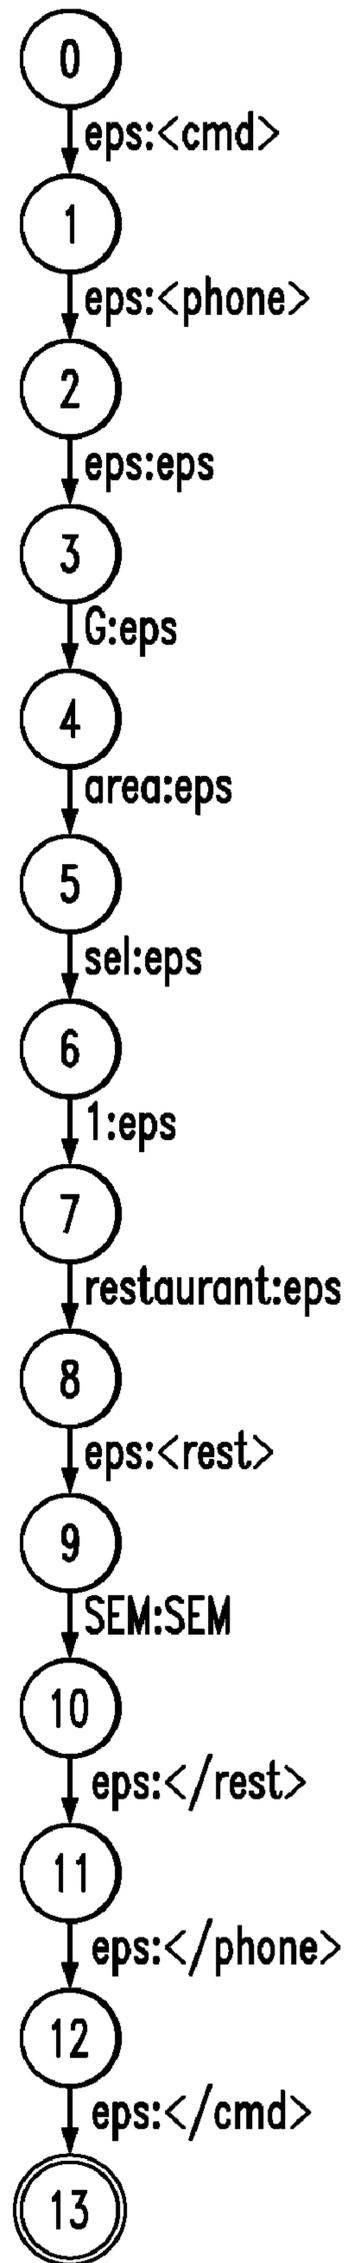


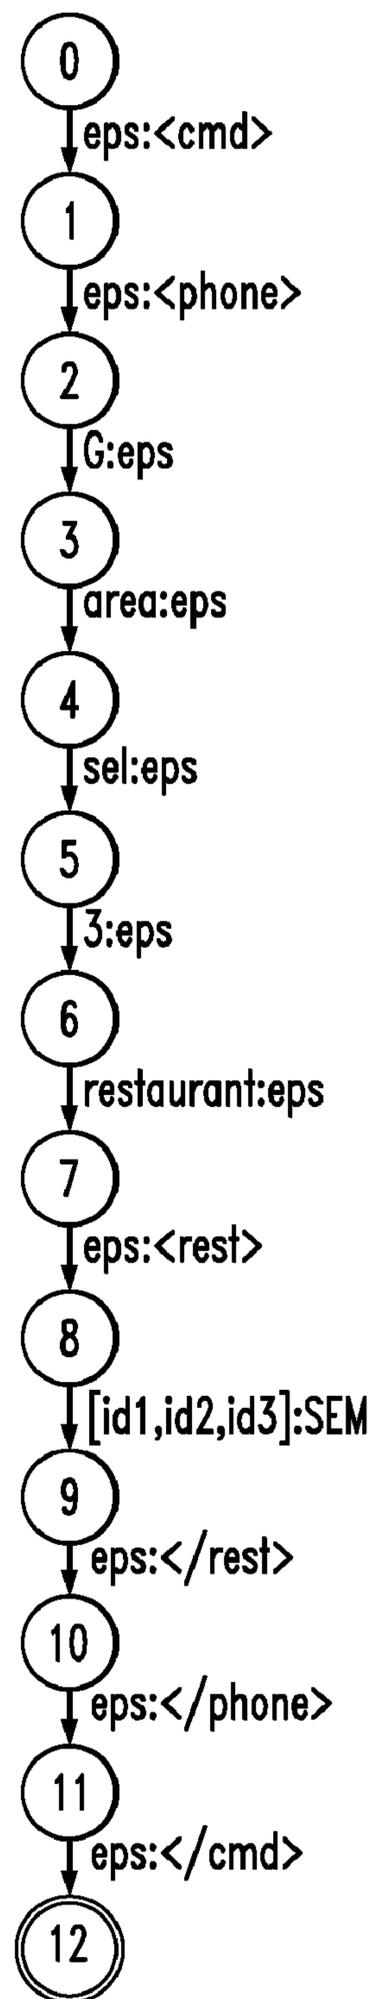
FIG. 22



*FIG. 23*



*FIG. 24*



*FIG. 25*

<cmd><phone><rest>[id1]</rest><rest>[id2,id3]</rest></phone></cmd>

**SYSTEMS AND METHODS FOR  
GENERATING MARKUP-LANGUAGE BASED  
EXPRESSIONS FROM MULTI-MODAL AND  
UNIMODAL INPUTS**

PRIORITY CLAIM

This application is a continuation of U.S. patent application Ser. No. 10/278,921, filed Oct. 24, 2002, the contents of which is incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

1. Field of Invention

This invention is directed to generating markup-language expressions from multimodal and/or unimodal inputs.

2. Description of Related Art

Multimodal interfaces allow input and/or output to be conveyed over multiple different channels, such as speech, graphics, gesture and the like. Multimodal interfaces enable more natural and effective interaction, because particular modes are best-suited for particular kinds of content. Multimodal interfaces are likely to play a critical role in the ongoing migration of interaction from desktop computing to wireless portable computing devices, such as personal digital assistants, like the Palm Pilot®, digital cellular telephones, public information kiosks that are wirelessly connected to the Internet or other distributed networks, and the like. One barrier to adopting such wireless portable computing devices is that they offer limited screen real estate, and often have limited keyboard interfaces, if any keyboard interface at all.

To realize the full potential of such wireless portable computing devices, multimodal interfaces need to support not just input from multiple modes. Rather, multimodal interfaces also need to support synergistic multimodal utterances that are optimally distributed over the various available modes. In order to achieve this, the content from different modes needs to be effectively integrated.

One previous attempt at integrating the content from the different modes is disclosed in "Unification-Based Multimodal Integration", M. Johnston et al., Proceedings of the 35th ACL, Madrid Spain, p. 281-288, 1997 (Johnston 1), incorporated herein by reference in its entirety. Johnston 1 disclosed a pen-based device that allows a variety of gesture utterances to be input through a gesture mode, while a variety of speech utterances can be input through a speech mode.

In Johnston 1, a unification operation over typed feature structures was used to model the integration between the gesture mode and the speech mode. Unification operations determine the consistency of two pieces of partial information. If the two pieces of partial information are determined to be consistent, the unification operation combines the two pieces of partial information into a single result. Unification operations were used to determine whether a given piece of gestural input received over the gesture mode was compatible with a given piece of spoken input received over the speech mode. If the gestural input was determined to be compatible with the spoken input, the two inputs were combined into a single result that could be further interpreted.

In Johnston 1, typed feature structures were used as a common meaning representation for both the gestural inputs and the spoken inputs. In Johnston 1, the multimodal integration was modeled as a cross-product unification of feature structures assigned to the speech and gestural inputs. While the technique disclosed in Johnston 1 overcomes many of the limitations of earlier multimodal systems, this technique does not scale well to support multi-gesture utterances, complex

unimodal gestures, or other modes and combinations of modes. To address these limitations, the unification-based multimodal integration technique disclosed in Johnston 1 was extended in "Unification-Based Multimodal Parsing", M. Johnston, Proceedings of COLING-ACL 98, p. 624-630, 1998 (Johnston 2), herein incorporated by reference in its entirety. The multimodal integration technique disclosed in Johnston 2 uses a multi-dimensional chart parser. In Johnston 2, elements of the multimodal input are treated as terminal edges by the parser. The multimodal input elements are combined together in accordance with a unification-based multimodal grammar. The unification-based multimodal parsing technique disclosed in Johnston 2 was further extended in "Multimodal Language Processing", M. Johnston, Proceedings of ICSLP 1998, 1998 (published on CD-ROM only) (Johnston 3), incorporated herein by reference in its entirety.

Johnston 2 and 3 disclosed how techniques from natural language processing can be adapted to support parsing and interpretation of utterances distributed over multiple modes. In the approach disclosed by Johnston 2 and 3, speech and gesture recognition produce n-best lists of recognition results. The n-best recognition results are assigned typed feature structure representations by speech interpretation and gesture interpretation components. The n-best lists of feature structures from the spoken inputs and the gestural inputs are passed to a multi-dimensional chart parser that uses a multimodal unification-based grammar to combine the representations assigned to the input elements. Possible multimodal interpretations are then ranked. The optimal interpretation is then passed on for execution.

SUMMARY OF THE INVENTION

An alternative to the unification-based multimodal parsing technique disclosed in Johnston 3 is discussed in "Finite-state Multimodal Parsing and Understanding", M. Johnston and S. Bangalore, Proceedings of COLING 2000, Saarbrücken, Germany, 2000, (Johnston 4) and in U.S. patent application Ser. No. 09/904,253, each incorporated herein by reference in its entirety. In Johnston 4, multimodal parsing, understanding and/or integration are achieved using a weighted finite-state device which takes speech and gesture streams as inputs and outputs their joint interpretation. This finite state approach is significantly more efficient, enables tight-coupling of multimodal understanding with speech recognition, and provides a general probabilistic framework for multimodal ambiguity resolution.

In Johnston 4 and the incorporated application, the symbols representing the multimodal inputs are concatenated together to obtain expressions that logically represent the multimodal inputs. For example, expressions such as, "email ([person(id1), organization(id2)])" are obtained. However, before the expressions can be used by underlying applications to perform functions based on the multimodal inputs, the expressions generally need to be converted into expressions that can be used by the other applications. Thus, a parser, which can parse the logical expressions, is used to parse or otherwise process the expressions before the expressions can be used to execute other functions. The parser must be able to parse the type of logical expressions which are created by concatenating the symbols representing the multimodal inputs and/or unimodal inputs.

This invention provides systems and methods for representing multimodal inputs and/or unimodal inputs using markup-language-based expressions.

This invention provides systems and methods for representing multimodal inputs and/or unimodal inputs using markup-language-based semantics.

This invention provides systems and methods for representing multimodal inputs and/or unimodal using extensible markup-language-based semantics.

This invention separately provides systems and methods for representing multimodal inputs and/or unimodal inputs using finite-state devices having markup-language-based symbols that represent the structure and the meaning of the multimodal and/or unimodal inputs.

This invention separately provides systems and methods for obtaining the meaning of multimodal inputs and/or by concatenating together the markup-language-based expressions representing the multimodal and/or unimodal inputs.

This invention separately provides systems and methods for obtaining the meaning of multimodal and/or unimodal inputs by concatenating together the extensible-markup-language-based expressions representing the multimodal and/or unimodal inputs.

This invention separately provides systems and methods for obtaining markup-language-based expressions representing multimodal and/or unimodal inputs using markup-language-based semantic representations of the multimodal inputs and/or unimodal inputs to perform desired functions.

This invention separately provides systems and methods for obtaining extended-markup-language-based expressions representing multimodal inputs and/or unimodal inputs using extended markup-language-based semantic representations of the multimodal inputs and/or unimodal inputs to perform desired functions.

This invention separately provides systems and methods for converting a sequence of symbols into a sequence of markup-language symbols.

This invention separately provides systems and methods for converting one markup-language expression into a second markup-language expression.

In various exemplary embodiments of the systems and methods according to this invention, information is input into a unimodal and/or multimodal device. The information is processed and/or converted using finite-state devices that contain markup-language based tags. The finite-state devices can represent regular grammars that contain markup-language tags in the terminals of the grammar. Markup-language based expressions, representing the inputted information and/or the processed information, are obtained by concatenating the symbols in the finite-state device representing the processed and/or converted information. The markup-language based expressions can be used other applications to further process and/or to store the information.

In various other exemplary embodiments of the systems and methods according to this invention, one or more commands are input into a unimodal and/or multimodal device. The commands are processed and/or converted using finite-state devices that contain markup-language based tags. The finite-state devices can represent regular grammars that contain markup-language tags in the terminals of the grammar. The finite-state devices can also represent approximated context-free grammars. Techniques for approximating context-free-grammars are described Fernando C. N. Pereira and R. Wright. 1997, "Finite-state approximation of phrase structure grammars," in E. Roche and Schabes Y., editors, Finite State Devices for Natural Language Processing. MIT Press, Cambridge, Mass. Markup-language based expressions, representing the inputted commands, are obtained by concatenating the symbols in the finite-state device representing the processed and/or converted commands. The markup-lan-

guage based expressions can be used other applications to further process and/or to execute the commands.

These and other features and advantages of this invention are described in, or are apparent from, the following detailed description of various exemplary embodiments of the systems and methods according to this invention.

#### BRIEF DESCRIPTION OF THE DRAWINGS

Various exemplary embodiments of this invention will be described in detail, with reference to the following figures, wherein:

FIG. 1 is a flowchart outlining one exemplary embodiment of a method for obtaining a markup-language expression from a plurality of multimodal inputs and/or unimodal inputs using finite-state devices;

FIG. 2 is a block diagram illustrating one exemplary embodiment of a conventional automatic speech recognition system usable with a multimodal meaning recognition system according to this invention;

FIG. 3 is a block diagram illustrating a first exemplary embodiment of a multimodal user input device and one exemplary embodiment of a multimodal meaning recognition system according to this invention;

FIG. 4 is a block diagram illustrating in greater detail one exemplary embodiment of the gesture recognition system of FIG. 3;

FIG. 5 is a block diagram illustrating in greater detail one exemplary embodiment of the multimodal parser and meaning recognition system of FIG. 3;

FIG. 6 is an exemplary embodiment of the multimodal user input device;

FIG. 7 is one exemplary embodiment of a three-tape multimodal grammar fragment usable by the multimodal meaning recognition system according to this invention;

FIG. 8 is one exemplary embodiment of a gesture finite-state machine generated by recognizing the gesture inputs shown in the exemplary embodiment of the multimodal user input device shown in FIG. 6;

FIG. 9 is an exemplary embodiment of an unimodal pen command inputted into the second exemplary embodiment of the multimodal user input device shown in FIG. 6;

FIG. 10 is one exemplary embodiment of two area gestures input into the multimodal user input device of FIG. 6;

FIG. 11 is one exemplary embodiment of a three-tape multimodal finite-state automaton usable to recognize the multimodal inputs received from the exemplary embodiment of the multimodal user input device shown in FIG. 6;

FIG. 12 is one exemplary embodiment of a multimodal grammar fragment usable by the multimodal recognition system according to this invention;

FIG. 13 is one exemplary embodiment of a gesture-to-speech finite-state transducer that represents the relationship between speech and gesture for the exemplary embodiment of the multimodal user input device shown in FIG. 6;

FIG. 14 is one exemplary embodiment of a speech/gesture/meaning finite-state transducer that represents the relationship between the combined speech and gesture symbols and the semantic meaning of the multimodal input for the exemplary embodiment of the multimodal input device shown in FIG. 6;

FIG. 15 is one exemplary embodiment of a gesture finite-state transducer generated by recognizing the gesture inputs shown in the exemplary embodiment of the multimodal user input device shown in FIG. 6;



## 5

FIG. 16 is one exemplary embodiment of a gesture finite-state machine generated by recognizing the gesture inputs shown in the exemplary embodiment of the multimodal user input device shown in FIG. 6;

FIG. 17 is a flowchart outlining one exemplary embodiment of a method for extracting meaning from a plurality of multimodal inputs;

FIG. 18 is one exemplary embodiment of a gesture/language finite-state transducer illustrating the composition of the gesture finite-state machine shown in FIG. 16 with the gesture-to-speech finite-state transducer shown in FIG. 13;

FIG. 19 is one exemplary embodiment of a speech input lattice generated by recognizing the speech input received by the exemplary embodiment of the multimodal user input device shown in FIG. 6;

FIG. 20 illustrates one exemplary embodiment of a gesture/speech finite-state transducer generated by composing the gesture/language finite-state transducer shown in FIG. 18 with the speech input lattice shown in FIG. 19;

FIG. 21 is one exemplary embodiment of a gesture/speech finite-state machine obtained from the gesture/speech finite-state transducer shown in FIG. 20;

FIG. 22 is one exemplary embodiment of a finite-state transducer, obtained from composing the gesture/speech finite-state machine shown in FIG. 20 with the speech/gesture/meaning finite-state transducer shown in FIG. 14, which extracts the meaning from the multimodal gestural and spoken inputs received when using the exemplary embodiment of the multimodal user input device shown in FIG. 6;

FIG. 23 is one exemplary embodiment of a finite-state machine generated by taking a projection of the finite-state transducer shown in FIG. 22;

FIG. 24 is one exemplary embodiment of a finite-state machine generated by composing the finite-state machine shown in FIG. 23 with the gesture finite-state transducer shown in FIG. 15; and

FIG. 25 is one exemplary illustration of a string of meaning read from the finite-state machine shown in FIG. 24.

#### DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

This invention provides systems and methods for representing multimodal inputs and/or unimodal inputs using markup-language-based expressions. Markup-language representation is considerably more expressive than the logical representation used in Johnston 4 and the incorporated application. XML is, for example, a widely-used markup language and thus, parsers for the XML are easier to obtain than parsers for the logical expressions generated by the systems and methods of used in Johnston 4. Further, because, for example, XML is a known standard, people are likely to be familiar with the expressions and thus, it is easier for a person to understand the expressions.

The extensible-markup language, XML is, for example, a known markup language for documents containing structured information, which is information that contains both content and some indication of the format of the content. Thus, the extensible-markup language can be used to store and/or to process the information contained in multimodal inputs and/or unimodal inputs of the systems and methods according to this invention. In various exemplary embodiments of the systems and methods according to this invention, extensible-markup-language semantics are used to represent the structured information in the multimodal inputs and/or unimodal inputs.

## 6

For example, the multimodal utterance “show Chinese restaurants in this neighborhood”, results in the following extensible-markup language representation:

---

```

<show>
  <restaurant>
    <cuisine> Chinese </cuisine>
    <location> ( . . . points . . . ) </location>
  </restaurant>
</show>.

```

---

Similarly, in another example, the multimodal utterance “show cheap Italian restaurants in Chelsea”, results in the following XML representation:

---

```

<cmd>
  <show>
    <rest>
      <cuisine> Italian </cuisine>
      <price> inexpensive </price>
      <location> Chelsea </location>
    </rest>
  </show>
</cmd>.

```

---

In exemplary embodiments of the systems and methods according to this invention, because markup-language tags and markup-language-based semantics are used in the multimodal grammar and the finite-state devices, a dedicated parser is not necessary to process the resulting expression. Various markup languages, such as XML and HTML, are known standards, and parsers are available for various markup languages, such as XML and HTML, in a broad range of programming languages. Thus, the obtained logical expressions need not be processed by a parser dedicated to parsing that type of logical expression before the expression can be used by underlying applications and other system components.

For example, in various exemplary embodiments of the systems and methods of this invention, upon receipt of the above extensible-markup language expression, for example, the extensible-markup language parser will parse the extensible-markup language expression. The extensible-markup language parser will determine that the expression represents a request to show a restaurant, for example. The parser will use the sub-attributes of the extensible-markup language tag <rest>, for example, to construct a database query. The query will look for the set of restaurants, for example, in the database, which meet the specified criteria. The resulting restaurants, for example, that meet the criteria are presented to the user.

The systems and methods according to this invention also facilitate logging of the system’s operation. The logging of the system’s operation can be used, for example, for user studies, evaluation, and system debugging. In various exemplary embodiments of the systems and methods according to this invention, not only can the meaning representation of the multimodal and/or unimodal input, for example, be represented with markup-language-based expressions, the communications between system components can also, for example, be encoded in markup-language-based expressions. For example, the communications between system components of the systems and methods according to this invention, can be recorded in a system log in extensible-markup-language format.

With the meaning representation and the communications between system components in extensible-markup-language format, for example, standard extensible-markup-language tools can, for example, be used to view the system's log. It is also possible to build custom log viewers which can use extensible-stylesheet-language transformations, i.e., XSLT and extensible-markup-language-path language, i.e., XPATH, for example, in order to access relevant components of the system log and render them as hypertext-markup language, i.e., HTML in a browser.

In addition, by creating extensible-markup-language-based expressions, for example, the systems and methods according to this invention can be easily mapped, for example, to typed feature structures. The mapping from extensible-markup-language representations to typed feature structures, for example, is disclosed in "The logic of typed feature structures", Bob Carpenter, Cambridge University Press, 1992. Typed features structures are commonly used in natural language processing and reasoning. By creating, for example, XML expressions representing the multimodal inputs, for example, the XML expressions can be mapped into typed feature structures. The typed feature structures can be used, for example, to determine how the incoming inputs relate to the prior dialogues.

Further, the document type description (DTD) mechanism of XML enables one to specify what is well formed for a particular subset of XML. In the systems and methods of this invention, well formed constraints can be specified in the meaning representation, for example. To debug the multimodal grammar, for example, all the possible meaning representations can be enumerated and each of the representations can be tested, for example, by the DTD. The testing can be used to determine, for example, the malformed meaning representations and/or the meaning representations which lie outside of the system's capabilities.

FIG. 1 is a flowchart outlining one exemplary embodiment of a method for generating markup-language-based expressions representing the meaning and/or a desired result of a performed function using data input into a multimodal and/or unimodal recognition device. Beginning in step 10, operation continues to step 20, where data is input into the multimodal and/or unimodal device. Then, in step 30, at least one input lattice is generated from the input data, where the input lattice represents the structure and the content of the input data. Next, in step 40, a function is performed using the generated input lattice and at least one other finite-state device. The at least one other finite-state device contains symbols based on the semantics of a markup language. Operation then continues to step 50.

In step 50, a result finite-state device is generated, where the result finite-state device represents the result of the performed function and contains symbols based on the semantics of the markup language. Next, in step 60, the symbols representing the result of the performed function are concatenated together to generate a markup language expression. Then, in step 70, the markup-language is output to another process. Operation then continues to step 80, where operation of the method ends.

It should be appreciated that various exemplary embodiments of the systems and methods according to this invention can be used to perform various functions using finite-state devices. In various exemplary embodiments, the systems and methods according to this invention can be used, for example, to perform a function of transforming a sequence of words to a sequence of markup-language symbols. It is possible, in various exemplary embodiments of the systems and methods according to this invention, for example, to perform a func-

tion of transforming a sequence of words to a sequence of markup-language symbols by using, for example, a finite-state transducer that, at least, maps words to markup-language symbols.

In various exemplary embodiment, the systems and methods according to this invention can be used to perform a function of converting one markup language expression into another markup-language expression. It is possible, for example, to perform a function of converting one markup-language expression into another markup-language expression by using a transducer which maps one markup-language expression to another.

As discussed below, in various exemplary embodiments, the systems and methods according to this invention can, for example, be used to perform a function of multimodal integration. Multimodal integration is performed, in various exemplary embodiments of the systems and methods of this invention, using a multimodal grammar that contains extensible-markup-language tags in, for example, the result parts of the terminals. Extensible-markup-language expressions can be obtained by concatenating together the expressions representing the results of the performed function. Because the result of multimodal integration is, for example, in the form of extensible-markup-language expressions, the resulting-concatenated expression, can, for example, be used by other devices and/or applications to further process the multimodal inputs. It should be appreciated, however, that, in various exemplary embodiments, the systems and methods according to this invention can be used to perform other applicable functions using finite-state devices.

FIG. 2 illustrates one exemplary embodiment of an automatic speech recognition system 100 usable with the multimodal recognition and/or meaning system 1000 according to this invention that is shown in FIG. 3. As shown in FIG. 2, automatic speech recognition can be viewed as a processing pipeline or cascade.

In each step of the processing cascade, one or two lattices are input and composed to produce an output lattice. In automatic speech recognition and in the following description of the exemplary embodiments of the systems and methods of this invention, the term "lattice" denotes a directed and labeled graph, which is possibly weighted. In each lattice, there is typically a designated start node "s" and a designated final node "t". Each possible pathway through the lattice from the start node s to the final node t induces a hypothesis based on the arc labels between each pair of nodes in the path. For example, in a word lattice, the arc labels are words and the various paths between the start node s and the final node t form sentences. The weights on the arcs on each path between the start node s and the final node t are combined to represent the likelihood that that path will represent a particular portion of the utterance.

As shown in FIG. 2, one exemplary embodiment of a known automatic speech recognition system 100 includes a signal processing subsystem 110, an acoustic model lattice 120, a phonetic recognition subsystem 130, a lexicon lattice 140, a word recognition subsystem 150, a grammar or language model lattice 160, and a task recognition subsystem 170. In operation, uttered speech is input via a microphone, which converts the sound waves of the uttered speech into an electronic speech signal. The electronic speech signal is input to the signal processing subsystem 110 on a speech signal input line 105. The signal processing subsystem 110 digitizes the electronic speech signal to generate a feature vector lattice 115. The feature vector lattice 115 is a lattice of acoustic feature vectors. The feature vector lattice 115 is input along with the acoustic model lattice 120 to the phonetic recogni-

tion subsystem **130**. The acoustic model lattice **120** represents a set of acoustic models and is applied to transform the feature vector lattice **115** into a phone lattice. Each node of the phone lattice represents a spoken sound, such as, for example, the vowel /e/ in “bed”.

The phone lattice **135** is input along with the lexicon lattice **140** into the word recognition subsystem **150**. The lexicon lattice **140** describes different pronunciations of various words and transforms the phone lattice **135** into a word lattice **155**. The word lattice **155** is then input, along with the grammar or language model lattice **160**, into the utterance recognition subsystem **170**. The grammar or language model lattice **160** represents task-specific information and is used to extract the most likely sequence of uttered words from the word lattice **155**. Thus, the utterance recognition subsystem **170** uses the grammar or language model lattice **160** to extract the most likely sentence or other type of utterance from the word lattice **155**. In general, the grammar or language model lattice **160** will be selected based on the task associated with the uttered speech. The most likely sequence of words, or the lattice of n most-likely sequences of words, is output as the recognized utterance **175**.

In particular, one conventional method of implementing automatic speech recognition forms each of the acoustic model lattice **120**, the lexicon lattice **140** and the grammar or language model lattice **160** as a finite-state transducer. Thus, each of the phonetic recognition subsystem **130**, the word recognition subsystem **150**, and the utterance recognition **170** performs a generalized composition operation between its input finite-state transducers. In addition, the signal processing subsystem **110** outputs the features vector lattice **115** as a finite-state transducer.

Conventionally, the grammar or language model lattice **160** is predetermined and incorporated into the automatic speech recognition system **100** based on the particular recognition task that the automatic speech recognition system **100** is to perform. In various exemplary embodiments, any of the acoustic model lattice **120**, the lexicon lattice **140** and/or the grammar or language model **160** can be non-deterministic finite-state transducers. In this case, these non-deterministic finite-state transducers can be determinized using the various techniques disclosed in “Finite-state transducers in Language and Speech Processing”, M. Mohri, Computational Linguistics, 23:2, p. 269-312, 1997, U.S. patent application Ser. No. 09/165,423, filed Oct. 2, 1998, and/or U.S. Pat. No. 6,073,098 to Buchsbaum et al., each incorporated herein by reference in its entirety.

In contrast, in various exemplary embodiments of the systems and methods according to this invention, in the multimodal recognition or meaning system **1000** shown in FIG. 3, the automatic speech recognition system **100** uses a grammar or language model lattice **160** that is obtained from the recognized gestural input received in parallel with the speech signal **105**. This is shown in greater detail in FIG. 3. In this way, the output of the gesture recognition system **200** can be used to compensate for uncertainties in the automatic speech recognition system.

Alternatively, in various exemplary embodiments of the systems and methods according this invention, the output of the automatic speech recognition system **100** and output of the gesture recognition system **200** can be combined only after each output is independently obtained. In this way, it becomes possible to extract meaning from the composition of two or more different input modes, such as the two different input modes of speech and gesture.

Furthermore, it should be appreciated that, in various exemplary embodiments of the systems and methods accord-

ing to this invention, the output of the gesture recognition system **200** can be used to provide compensation to the automatic speech recognition system **100**. Additionally, their combined output can be further processed to extract meaning from the combination of the two different input modes. In general, when there are two or more different input modes, any of one or more of the input modes can be used to provide compensation to one or more other ones of the input modes.

Thus, it should further be appreciated that, while the following detailed description focuses on speech and gesture as the two input modes, any two or more input modes that can provide compensation between the modes, which can be combined to allow meaning to be extracted from the two or more recognized outputs, or both, can be used in place of, or in addition to, the speech and gesture input modes discussed herein.

In particular, as shown in FIG. 3, when speech and gesture are the implemented input modes, a multimodal user input device **400** includes a gesture input portion **410** and a speech input portion **420**. The gesture input portion **410** outputs a gesture signal **205** to a gesture recognition system **200** of the multimodal recognition and/or meaning system **1000**. At the same time, the speech input portion **420** outputs the speech signal **105** to the automatic speech recognition system **100**. The gesture recognition system **200** generates a gesture recognition lattice **255** based on the input gesture signal **205** and outputs the gesture recognition lattice **255** to a multimodal parser and meaning recognition system **300** of the multimodal recognition and/or meaning system **1000**.

In those various exemplary embodiments that provide compensation between the gesture and speech recognition systems **200** and **100**, the multimodal parser/meaning recognition system **300** generates a new grammar or language model lattice **160** for the utterance recognition subsystem **170** of the automatic speech recognition system **100** from the gesture recognition lattice **255**. In particular, this new grammar or language model lattice **160** generated by the multimodal parser/meaning recognition system **300** is specific to the particular sets of gestural inputs generated by a user through the gesture input portion **410** of the multimodal user input device **400**. Thus, this new grammar or language model lattice **160** represents all of the possible spoken strings that can successfully combine with the particular sequence of gestures input by the user through the gesture input portion **410**. That is, the recognition performed by the automatic speech recognition system **100** can be improved because the particular grammar or language model lattice **160** being used to recognize that spoken utterance is highly specific to the particular sequence of gestures made by the user.

The automatic speech recognition system **100** then outputs the recognized possible word sequence lattice **175** back to the multimodal parser/meaning recognition system **300**. In those various exemplary embodiments that do not extract meaning from the combination of the recognized gesture and the recognized speech, the recognized possible word sequences lattice **175** is then output to a downstream processing task. The multimodal recognition and/or meaning system **1000** then waits for the next set of inputs from the multimodal user input device **400**.

In contrast, in those exemplary embodiments that additionally extract meaning from the combination of the recognized gesture and the recognized speech, the multimodal parser/meaning recognition system **300** extracts meaning from the combination of the gesture recognition lattice **255** and the recognized possible word sequences lattice **175**. Because the spoken utterances input by the user through the speech input portion **420** are presumably closely related to the gestures

input at the same time by the user through the gesture input portion 410, the meaning of those gestures can be tightly integrated with the meaning of the spoken input generated by the user through the speech input portion 420.

The multimodal parser/meaning recognition system 300 outputs a recognized possible meaning lattice 375 in addition to, or in place of, one or both of the gesture recognition lattice 255 and/or the recognized possible word sequences lattice 175. In various exemplary embodiments, the multimodal parser and meaning recognition system 300 combines the recognized lattice of possible word sequences 175 generated by the automatic speech recognition system 100 with the gesture recognition lattice 255 output by the gesture recognition system 200 to generate the lattice of possible meaning sequences 375 corresponding to the multimodal gesture and speech inputs received from the user through the multimodal user input device 400.

Moreover, in contrast to both of the embodiments outlined above, in those exemplary embodiments that only extract meaning from the combination of the recognized multimodal inputs, the multimodal parser/meaning recognition system 300 does not generate the new grammar or language model lattice 160. Thus, the gesture recognition lattice 255 does not provide compensation to the automatic speech recognition system 100. Rather, the multimodal parser/meaning recognition system 300 only combines the gesture recognition lattice 255 and the recognized possible word sequences lattice 175 to generate the recognition meaning lattice 375.

When the gesture recognition system 200 generates only a single recognized possible sequence of gestures as the gesture recognition lattice 255 that means there is essentially no uncertainty in the gesture recognition. In this case, the gesture recognition lattice 255 provides compensation to the automatic speech recognition system 100 for any uncertainty in the speech recognition process. However, the gesture recognition system 200 can generate a lattice of n possible recognized gesture sequences as the gesture recognition lattice 255. This recognizes that there may also be uncertainty in the gesture recognition process.

In this case, the gesture recognition lattice 255 and the word lattice 155 provide mutual compensation for the uncertainties in both the speech recognition process and the gesture recognition process. That is, in the face of this uncertainty, the best, i.e., most-probable, combination of one of the n-best word sequences in the word lattice 155 with one of the n-best gesture sequences in the gesture recognition lattice may not include the best recognition possible sequence from either the word lattice 155 or the gesture recognition lattice 255. For example, the most-probable sequence of gestures in the gesture recognition lattice may combine only with a rather low-probability word sequence through the word lattice, while the most-probable word sequence may combine well only with a rather low-probability gesture sequence. In contrast, a medium-probability word sequence may match very well with a medium-probability gesture sequence. Thus, the net probability of this latter combination of word and gesture sequences may be higher than the probability of the combination of the best word sequence with any of the gesture sequences through the gesture recognition lattice 255 and may be higher than the probability of the combination of the best gesture sequence with any of the word sequences through the lattice of possible word sequences 155. In this way, mutual compensation is provided between the gesture recognition system 200 and the automatic speech recognition system 100.

FIGS. 4 and 5 illustrate in greater detail various exemplary embodiments of the gesture recognition system 200, the mul-

timodal parser/meaning recognition system 300, and the multimodal user input device 400. In particular, as shown in FIG. 4, one exemplary embodiment of the gesture recognition system 200 includes a gesture feature extraction subsystem 210 and a gesture recognition subsystem 230. Various other exemplary embodiments may include a gesture language model lattice and a gesture meaning subsystem. In operation, gesture utterances are input through the gesture input portion 410 of the multimodal user input device 400, which converts the movements of an input device, such as a mouse, a pen, a trackball, a track pad or any other known or later-developed gestural input device, into an electronic gesture signal 205. At the same time, the multimodal user input device 400 converts the gestural input into digital ink that can be viewed and understood by the user. This is shown in greater detail in FIG. 6.

The gesture feature extraction subsystem 210 converts the motions of the gesture input device represented by the gesture signal 205 into a gesture feature lattice 220. As disclosed in Johnston 1-3, the various gestures that can be made can be as simple as pointing gestures to a particular information element at a particular location within the gesture input portion 410 of the multimodal user input device 400, or can be as complex as a specialized symbol that represents a type of military unit on a military map displayed in the gesture input portion 410 of the multimodal user input device 400 and includes an indication of how the unit is to move, and which unit is to move and how far that unit is to move, as described in detail in Johnston 1.

The gesture feature lattice 220 is input to the gesture recognition subsystem 230. The gesture recognition subsystem 230 may be implemented as a neural network, as a Hidden-Markov Model (HMM) or as a simpler template-based gesture classification algorithm. The gesture recognition subsystem 230 converts the gesture feature lattice 220 into the gesture recognition lattice 255. The gesture recognition lattice 255 includes the identities of graphical elements against which deictic and other simple "identification" gestures are made, possible recognition of more complex gestures that the user may have made and possibly the locations on the displayed graphics where the more complex gesture was made, such as in Johnston 1, and the like. As shown in FIG. 3, the gesture recognition system 200 outputs the gesture recognition lattice 255 to the multimodal parser/meaning recognition system 300.

It should be appreciated that the gesture feature recognition subsystem 210 and the gesture recognition subsystem 230 can each be implemented using any known or later-developed system, circuit or technique that is appropriate. In general, the entire gesture recognition system 200 can be implemented using any known or later-developed system that generates a directed graph from a gesture input. For example, for any system or device embodiment, a computer readable medium can store instructions for controlling the computing device to perform steps for any system or device embodiment, a computer readable medium or memory can store instructions for controlling a computing device via its processor to perform steps disclosed herein. Such computer readable medium excludes propagation media such as signals per se and carrier waves.

For example, one known system captures the time and location or locations of the gesture. Optionally, these inputs are then normalized and/or rotated. The gestures are then provided to a pattern classification device that is implemented as part of the gesture feature recognition subsystem 210. In various exemplary embodiments, this pattern classification device is a template matching system, which transforms the

gesture into a feature vector. In various other exemplary embodiments, this pattern classification device is a neural network or a Hidden Markov Model that has been trained to recognize certain patterns of one or more temporally and/or spatially related gesture components as a specific set of features.

When a single gesture is formed by two or more temporally and/or spatially related gesture components, those gesture components can be combined into a single gesture either during the recognition process or by the multimodal parser/meaning recognition system 300. Once the gesture features are extracted, the gesture recognition subsystem 230 combines the temporally adjacent gestures into a lattice of one or more recognized possible gesture sequences that represent how the recognized gestures follow each other in time.

In various exemplary embodiments, the multimodal parser and meaning recognition system 300 can be implemented using a single three-tape finite-state device that inputs the output lattices from the speech recognition system 100 and the gesture recognition system 200 and directly obtains and outputs a meaning result. In various exemplary embodiments, the three-tape finite-state device is a three-tape grammar model that relates the gestures and the words to a meaning of the combination of a gesture and a word. FIG. 7 shows a portion of such a three-tape grammar model usable in the multimodal parser and meaning recognition system 300 to generate a meaning output from gesture and speech recognition inputs. In general, the multimodal parser and meaning recognition system 300 can be implemented using an n-tape finite-state device that inputs n-1 lattices from a plurality of recognition systems usable to recognize an utterance having a plurality of different modes.

FIG. 5 shows the multimodal parser/meaning recognition system 300 in greater detail. As shown in FIG. 5, the multimodal parser/meaning recognition system 300 may include one or more of a gesture-to-speech composing subsystem 310, a gesture-to-speech finite-state transducer 320, a lattice projection subsystem 330, a gesture and speech composing subsystem 340, a speech/gesture combining subsystem 350, a speech/gesture/meaning lattice 360 and/or a meaning recognition subsystem 370. In particular, the gesture-to-speech composing subsystem 310 inputs the gesture recognition lattice 255 output by the gesture recognition system 200 and composes it with the gesture-to-speech finite-state transducer 320 to generate a gesture/language finite-state transducer 325. The gesture/language finite-state transducer 325 is output to both the lattice projection subsystem 330 and the gesture and speech composing subsystem 340.

The lattice projection subsystem 330 generates a projection of the gesture/language finite-state transducer 325 and outputs the projection of the gesture/language finite-state transducer 325 as the grammar or language model lattice 160 to the automatic speech recognition system 100. Thus, if the multimodal parser/meaning recognition system 300 does not also extract meaning, the gesture and speech composing subsystem 340, the speech/gesture combining subsystem 350, the speech/gesture/meaning lattice 360 and the meaning recognition subsystem 370 can be omitted. Similarly, if the multimodal parser/meaning recognition system 300 does not generate a new grammar or language model lattice 160 for the automatic speech recognition system 100, at least the lattice projection subsystem 330 can be omitted.

In those various embodiments that combine the gesture recognition lattice 255 and the recognized possible lattice of word sequences 175, whether or not the automatic speech recognition 100 has generated the lattice of possible word sequences 175 based on using the projection of the gesture/

language finite-state transducer 325 as the grammar or language model or lattice 160, the lattice of possible word sequences 175 is input by the multimodal parser/meaning recognition system 300. In particular, the gesture and speech composing subsystem 340 inputs both the lattice of possible word sequences 175 and the gesture/language finite-state transducer 325. In those various exemplary embodiments that do not use the output of the gesture recognition system 200 to provide compensation between the speech and gesture recognition systems 100 and 200, the gesture/language finite-state transducer 325 can be generated using any known or later-developed technique for relating the gesture recognition lattice 255 to the recognized possible lattice of word sequences 175 in place of the gesture-to-speech composing subsystem 310 and the gesture-to-speech finite-state transducer 320.

In those various exemplary embodiments that extract meaning from the multimodal inputs, the gesture and speech composing subsystem 340 composes these lattices to generate a gesture/speech finite-state transducer 345. The gesture and speech composing subsystem 340 outputs the gesture/speech finite-state transducer 345 to the speech/gesture combining subsystem 350. The speech/gesture combining subsystem 350 converts the gesture/speech finite-state transducer 345 to a gesture/speech finite-state machine 355. The gesture/speech finite-state machine 355 is output by the speech/gesture combining subsystem 350 to the meaning recognition subsystem 370. The meaning recognition subsystem 370 composes the gesture/speech finite-state machine 355 with the speech/gesture/meaning finite-state transducer 360 to generate a meaning lattice 375. The meaning lattice 375 combines the recognition of the speech utterance input through the speech input portion 420 and the recognition of the gestures input through the gesture input portion 410 of the multimodal user input device 400. The most probable meaning is then selected from the meaning lattice 375 and output to a downstream task.

To illustrate one exemplary embodiment of the operation of the multimodal recognition and/or meaning system 1000, the multimodal user input device 400 includes, for example, the gesture input portions 410 and the speech input portion 420 shown in FIG. 6. In particular, the gesture input portion 410 displays a graphical user interface that allows the user to direct either e-mail messages or pager messages to the various persons, departments, and/or organizations represented by the objects 412 displayed in the gesture input portion 410. The multimodal user input device 400 also allows the user to input spoken commands to the speech input portion, or microphone, 420.

For example, in this exemplary embodiment, the user interacts with a graphical interface displaying restaurant listings and a dynamic map showing locations and street information. In various exemplary embodiments, the gesture input portion 410 can be used to display, for example, a graphical user interface that includes a working-city-guide application and navigation screen that, for example, enables mobile users to access restaurant and subway information for a city. In this example, the user is able to interact with the graphical interface to display restaurant listings and a dynamic map showing locations and street information. This graphical user interface responds to commands entered using pen-based gestures and/or spoken utterances.

The user is, for example, free to give commands or to reply to requests displayed on the graphical user interface using speech, by drawing on the display with a stylus, or using synchronous multimodal combinations of the available modes. The user can, for example, ask for the review, cuisine,

phone number, address, or other information for a restaurant or set of restaurants. The working-city-guide application, for example, generates graphical callouts on the display.

For example, a user can request to see restaurants using the spoken command "Show cheap Italian restaurants in Chelsea". The working-city-guide application will then display the appropriate map location and show the locations of the restaurants that meet the indicated criteria. Alternatively, the user can give the same command multimodally by circling an area on the map and uttering "Show cheap Italian restaurants in this neighborhood". As shown in FIG. 9, if the immediate environment is too noisy or public, the same command can be input solely using a writing gesture, by circling an area of the displayed map and by writing "cheap" and "Italian". As shown in FIG. 10, for example, the user can speak the utterance "Phone numbers for these restaurants" and circle, for example, a total of three restaurants.

The working-city-guide application can also provide sub-way directions. For example, if the user speaks the spoken utterance "How do I get to this place?" and circles one of the restaurants displayed on the map, the working-city-guide application will, for example, display a query or graphical callout, such as, for example, the text string "Where do you want to go from?" The user can then respond, for example, with a spoken utterance stating a location, by saying a location, such as, for example, "25th Street and 3rd Avenue", with a gesture defining a location, for example, by writing "25th Street and 3rd Avenue", or by pointing to 25th Street and 3rd Avenue on the map, or multimodally by speaking the spoken utterance "from here" and pointing to or circling the location on the map.

The working-city-guide application then determines a sub-way route. In various exemplary embodiments, as appropriate, the working city-guide-application generates a multimodal presentation indicating the series of actions the user needs to take will be displayed. In various exemplary embodiments, the working-city-guide application starts, for example, by zooming in on the first station and then gradually presenting each stage of the route along with a series of synchronized test-to-speech (TTS) prompts.

The structure interpretation of multimodal commands of this kind can be captured declaratively in a multimodal regular grammar. A multimodal regular grammar can be defined formally as the quadruple MCFG as follows:

TABLE 1

---

MCFG =  $\langle N, T, P, S \rangle$  where  
 N is the set of non-terminals;  
 P is the set of projections of the form:  
 $A \rightarrow \alpha$  where  $A \in N$  and  $\alpha \in (N \cup T)^*$ ;  
 S is the start symbol for the grammar;  
 T is the set of terminals:  
 $((W \cup \epsilon) \times (G \cup \epsilon) \times (M \cup \epsilon))^+$ ,  
 where W is the vocabulary of the speech; G is the vocabulary of gesture;  
 $G = (\text{GestureSymbols} \cup \text{EventSymbols})$ ;  $\text{GestureSymbols} = \{\text{Gp}, \text{Go}, \text{Gpf}, \text{Gpm} \dots\}$ ;  
 Finite collections of  $\text{EventSymbols} = \{e_1 e_2 \dots\}$ ; and  
 M is the vocabulary that represents meaning and includes  $\text{EventSymbols} \subset M$ .

---

In general, a regular grammar can be approximated by a finite-state automaton. The transition symbols of the finite-state automaton are the terminals of the regular grammar. In the case of the multimodal regular grammar defined above, these terminals contain three components, W, G and M. With respect to the discussion outlined above regarding temporal constraints, more specific temporal constraints than order can be encoded in the finite-state approach by writing symbols

representing the passage of time onto the gesture tape and referring to such symbols in the multimodal grammar.

In the exemplary embodiment of the gesture input portion 410 shown in FIG. 6, the gestures are simple deictic circling gestures. The gesture meaning subsystem 250 assigns semantic types to each gesture based on the underlying portion of the gesture input portion 410 against which the gestures are made. In the exemplary embodiment shown in FIG. 6, the gestures are made relative to the objects 412 that can represent people, organizations or departments to which an e-mail message or a pager message can be directed. If the gesture input portion 410 were instead a map, the gestures would be referenced against particular map coordinates, where the gesture indicates the action to be taken at particular map coordinates or the location of people or things at the indicated map location.

In Johnston 4 and the incorporated 253 application compared with a feature-based multimodal grammar, these semantic types constitute a set of atomic categories which make the relevant distinctions for gesture events to predict speech events and vice versa. For example, if the gesture is a deictic, i.e., pointing, gesture to an object in the gesture input portion 410 that represents a particular person, then spoken utterances like "this person", "him", "her", and the like, are the preferred or predicted speech events and vice versa. These categories also play a role in constraining the semantic representation when the speech is underspecified with respect to the semantic type, such as, for example, spoken utterances like "this one".

In Johnston 4 and the incorporated 253 application, the gesture symbols G can be organized into a type hierarchy reflecting the ontology of the entities in the application domain. For example, in the exemplary embodiment of the gesture input portion 410 shown in FIG. 5, a pointing gesture may be assigned the general semantic type "G". This general semantic gesture "G" may have various subtypes, such as "Go" and "Gp", where "Go" represents a gesture made against an organization object, while the "Gp" gesture is made against a person object. Furthermore, the "Gp" type gesture may itself have subtypes, such as, for example, "Gpm" and "Gpf" for objects that respectively represent male and female persons.

The systems and methods for recognizing and representing gestures disclosed in U.S. patent application Ser. No. 10/216,390, incorporated herein by reference in its entirety, provide an approach that can be used instead of, or in addition, to using atomic symbols in the multimodal grammar and the corresponding finite-state machines to represent different types of gestures, gesture interpretations are encoded as sequences of symbols. In this approach, each symbol conveys a specific attribute of the content of the gesture, such as type or number. Under this approach, the Gesture Symbols and Event Symbols in the multimodal grammar fragment, shown in Table 1 are instead defined as:

TABLE 2

---

$\text{GestureSymbols} = \{G, \text{area}, \text{location}, \text{restaurant}, 1, \dots\}$ ;  
 $\text{EventSymbols} = \{\text{SEM}\}$

---

These definitions can be used instead of, or in addition to, the definitions of the Gesture Symbols and Event Symbols shown in Table 1.

By decomposing the gesture symbols into sequences of symbols, it is easier to reference sets of entities of a specific type. In addition, a smaller number of symbols are required in the alphabet of symbols that represent the gestural content of

the grammar. Further, decomposing the gesture symbols into sequences of symbols facilitates storing specific gesture content, discussed below, and aggregating adjacent selection gestures, as disclosed in U.S. patent application Ser. No. 10/216,448.

Under this approach, for example, the gesture symbol complexes have a basic format such as:

#### G Form Meaning (Number Type) SEM

However, it should be appreciated that the gesture symbol complexes can be implemented in any appropriate format. The "FORM" term specifies the physical form of the gesture. In various exemplary embodiments, the "FORM" term can take such values as area, point, line and arrow. The "MEANING" term specifies the specific meaning of the form of the gesture. For example, if the value for the "FORM" term of a gesture is area, the value for the "MEANING" term of that "FORM" term can be location, selection or any other value that is appropriate. If the value of the "MEANING" term of that "FORM" term is selection, such that one or more specific entities are selected, the "NUMBER" term and the "TYPE" term can be used to further specify the entities that are or have been selected. In particular, the value for the "NUMBER" term specifies the number of entities selected, such as 1, 2, 3, "many" and the like. Similarly, the value for the "TYPE" term specifies a particular type of entity, such as restaurant, theatre and the like, as appropriate for the given implementation or use. A value of mixed can be used for the "TYPE" term when one or more associated gestures reference a collection of entities of different types. The "SEM" term is a place holder for the specific content of the gesture, such as the points that make up an area or the identifiers (ids) of objects in a selection. To facilitate recomposing specific gestural content, specific content is mapped to a distinguished symbol, such as, the "SEM" term, while the other attributes of the gesture are mapped to themselves.

When using finite-state methods, in order to capture multimodal integration, it is necessary to abstract over specific aspects of gestural content. In the systems and methods disclosed in U.S. patent application Ser. No. 10/216,392, incorporated herein by reference in its entirety, abstraction is performed by representing the input gesture as a finite-state transducer that maps the specific contents of the input gesture to a distinguished symbol, such as, for example, the "SEM" term, while the other attributes of the input gesture are mapped to themselves. To perform multimodal integration, the output projection of the gesture finite-state transducer is used. After multimodal integration is completed, a projection of the gesture input and the meaning portions of the resulting finite-state machine is taken. The projection of the resulting finite-state machine is composed with the gesture input finite-state transducer to reintegrate the specific content of the gesture input which was left out of the finite-state process.

Thus, in the finite-state approach used in the systems and methods disclosed in incorporated Ser. No. 10/216,392 application, the specific content of the input gesture is essentially stored in the gesture finite-state transducer and a projection of the output of the gestural input finite-state transducer is used to conduct multimodal modal integration using finite-state devices. Multimodal integration is performed and a resulting finite-state device, which relates the gesture input and at least one other mode of multimodal input to meaning, is generated. After multimodal integration is performed, a projection of the resulting finite-state device is taken such that the projection contains the gesture input and the meaning portions of the resulting finite-state device. The specific content of the input

gesture is then retrieved from the gesture finite-state transducer by composing the gesture finite-state transducer with the projection of the resulting finite-state device.

In addition, in one exemplary embodiment of the systems and methods disclosed in the incorporated Ser. No. 10/216,392 application, a projection, i.e., the gesture input finite-state machine, is then used to perform multimodal integration. Accordingly, in order to reintegrate the specific content of the input gesture after multimodal integration is performed, an output projection of the resulting meaning finite-state machine is composed with the gesture finite-state transducer.

Further, in a multimodal system that integrates gestures and speech, deictic numeral expressions, such as, for example, "these four restaurants" can refer to a multitude of different possible sequences of gestures that are compatible and should be integrated with the spoken phrase. For example, the user might circle four individual restaurants, a set of four restaurants, two sets of two restaurants or a set of three restaurants and a single restaurant, etc. Therefore, in a multimodal recognition system that integrates gestures and speech, for example, capturing all of the gestural possibilities in the spoken language grammar significantly increases the size and complexity of the finite state devices used to represent the multimodal recognition. Thus, for such a known multimodal recognition system, every time a numerical expression exists in a multimodal input, a multitude of the alternative multimodal subcategorization frames would need to be assigned to that numerical expression by that known multimodal recognition system. Thus, any plural expression can be massively ambiguous.

This problem is presented in "Deixis and Conjunction in Multimodal Systems", M. Johnston, Proceedings of COLING 2000, Saarbruecken, Germany, 2000 (Johnston 5), herein incorporated by reference in its entirety. In Johnston 5, the problem is addressed by under specifying the configuration of gestures in the multimodal sub-categorization of the deictic numeral expressions. For example, instead of subcategorizing for a sequence of N gestures, the deictic numeral expression, 'these N', for example, subcategorizes for a collection of plurality N: <G[number:N]>. For example, the expression 'these four people' is subcategorized into <Gperson[number:4]>. Gesture combination rules enumerate the different ways to refer to a collection of entities.

In the systems and methods disclosed in Johnston 4 and the incorporated 253 application, certain simplifying assumptions are used with respect to temporal constraints in a multimodal recognition system that integrates gestures and speech. For multi-gesture utterances, the primary function of temporal constraints is to force an order on the gestures. For example, if a user generates the spoken utterance "move this here" and simultaneously makes two gestures, then the first gesture corresponds to the spoken utterance "this", while the second gesture corresponds to the spoken utterance "here".

In the various exemplary embodiments of the systems and methods according to the invention described in Johnston 4 and the incorporated 253 application, the multimodal grammars encode order, but do not impose explicit temporal constraints. In Johnston 4 and the incorporated 253 application, the deictic numerals are subcategorized for a plurality of the appropriate number and a set of gesture combination rules are predicatively applied to combine the elements of the gestural input into the appropriate pluralities.

In the systems and methods disclosed in U.S. patent application Ser. No. 10/216,448, incorporated herein by reference in its entirety, aggregation is used to preprocess the input lattice. For example, "gesture aggregation" is used to preprocesses the gesture lattice. Aggregation adds one or more

additional branches that represent combinations of related inputs to the lattice. For example, the inputs may be related spatially and/or temporally and can represent entities of the same or different type. The aggregated branch can then be combined with another input mode of the multimodal device, such as, for example, speech input that refers to the combination of the input gestures represented by the aggregated branch.

Aggregating temporally-related and/or spatially-related inputs according to the systems and methods of this invention significantly simplifies the construction of the multimodal grammar. The multimodal grammar can be significantly simplified by reducing, and ideally eliminating, the need to include all the possible combinations that could be assigned to such inputs, such as, for example, deictic numeral expressions. Accordingly, in various exemplary embodiments, the systems and methods of the incorporated Ser. No. 10/216,448 application also allow the size the finite state devices which are used to represent the multimodal grammar to be reduced, sometimes significantly. For example, in various exemplary embodiments, the finite state devices which are used for multimodal integration will often be significantly smaller than they otherwise would have been using this known technique.

For example, using the gesture symbol representation scheme disclosed in U.S. patent application Ser. No. 10/216,390, incorporated herein by reference in its entirety, an expression such as “these three restaurants” is assigned, for example, the gesture stream “G area selection 3 restaurant SEM” by the multimodal grammar. The spoken utterance stream will combine with a single area gesture containing three restaurants. However, to allow this single spoken utterance to readily combine with a sequence of three temporally-related but spatially-separated gestures, for example, aggregation should first be applied. In this example, aggregation is used to preprocess the gesture lattice by parsing the gesture lattice and composing together temporally-adjacent selections of identical type. Additional branches, representing the composed selections, are added to the gesture lattice.

For example, if the user makes three gestures, such as, “G area selection 1 restaurant SEM, G area selection 1 restaurant SEM and G area selection 1 restaurant SEM”, the resulting gesture lattice could include an aggregation of all three gestures and/or an aggregation of each temporally-related and/or spatially-related pair.

In various exemplary embodiments of the systems and methods disclosed in the incorporated Ser. No. 10/216,448 application, it is also possible to maintain a record of the original inputs, i.e., a record of the non-aggregated inputs. A record of the original sequence of inputs can be preserved, for example, by assigning weights to the paths in the resulting aggregated lattice. For example, in various exemplary embodiments, the newly aggregated paths can be assigned lower weights than the paths representing the original input sequences. Alternatively, for example, in various exemplary embodiments, a record of the original sequence of inputs can be maintained by adding a symbol, such as, for example, the symbol ‘AGG’, to the multimodal grammar and the resulting aggregated paths in the finite-state device. In addition, in various exemplary embodiments of the systems and methods according to this invention, the grammar may also be set up to aggregate only certain expressions.

In the finite-state automata approach used in the systems and methods according to the incorporated Ser. No. 10/216,448 application, in addition to capturing the structure of language with the finite-state device, meaning is also captured. This is significant in multimodal language processing, because the central goal is to capture how the multiple modes

contribute to the combined interpretation. In the finite-state automata technique used in the systems and methods according to this invention, symbols are written onto the third tape of the three-tape finite-state automaton, which, when concatenated together, yield the semantic representation for the multimodal utterance.

As discussed above, the user can speak the utterance “Phone numbers for these three restaurants” and circle, for example, as shown in FIG. 10, a total of three restaurants. As shown in FIG. 10, the user may combine the deictic spoken expression, “these three restaurants”, for example, with two area gestures, one which selects one restaurant and one which selects the other two restaurants. Alternatively, the user may, for example, combine the deictic spoken expression, “these three restaurants” with three separate selection gestures.

In a multimodal recognition system that integrates gestures and speech, to combine a deictic spoken utterance with multiple gestural inputs according to the systems and methods of the incorporated Ser. No. 10/216,448 application, the gesture lattice is pre-processed and spatially-related and/or temporally-related gestures are composed. The related gestures may, for example, be selections of identical type, such as, for example, a plurality of restaurant selections. Alternatively, depending on the spoken utterance, for example, the related gestures may, for example, be selections of mixed type, such as, for example, a restaurant and a theatre selection. Based on the compositions, branches are added to the gesture lattice.

It should be appreciated that the systems and methods according to this invention apply aggregation to combine gestures or inputs that are of identical type, i.e., type-specific aggregation. That is, type-specific aggregation refers to combining multiple gestures, such as for example, multiple gestures referring only to restaurants or multiple gestures referring only to theatres.

However, the systems and methods of this invention may also be used to combine gestures or inputs of mixed type, i.e., non-type-specific aggregation. That is, non-type-specific aggregation refers to combining multiple gestures, such as, for example, a gesture referring to one type of entity with another gesture referring to a second type of entity. For example, one gesture may refer to a restaurant, while a second gesture may refer to a theatre. For example, if the user says “tell me about these two” and circles a restaurant and a theatre, non-type-specific aggregation applies to combine the two gestures into an aggregate gesture or path of mixed type. For example, an aggregated path that can combine with the speech “these two”, can be: “G: G area: area selection: selection 2:2 mixed: mixed ([id4,id5]):SEM”.

In addition, the systems and methods according to this invention also include systems and methods for preserving an indication of the original sequences of the gestures. For example, in various exemplary embodiments, it is possible to preserve the original sequence of the gestures by assigning lower weights to the paths that were added as a result of aggregation.

Further, in order to abstract over specific aspects of the gestural content in accordance with the systems and methods of this invention, the gesture lattice is converted to a gesture transducer I•G, where the G side is the set of gesture symbols (including SEM) and I contains both gesture symbols and the specific contents. The specific contents of the gestural input includes, for example, entities or points on the gesture input portion 410 that are selected by the user. Further, the specific content may, for example, be an identification number, i.e., an entity identifier, that represents the entity or point on the gesture input portion. For example, to convert a gesture lattice



to a gesture transducer for example, a path reciting SEM([id1, id2, id3]) becomes [id1, id2, id3]: SEM.

Additionally, a gesture transducer I:G may be generated by the gesture recognition system based on the gestural input. In this case, the gesture transducer is converted to a gesture  
5 finite-state machine usable to carry out multimodal integration or any other applicable function using finite-state devices. For example, if the user circles two restaurants and the user says “phone numbers for these two restaurants” the gesture is represented as a transducer, where the transducer  
10 I:G has same symbol on each side except for the SEM arcs, which are split. In this case, to carry out the multimodal integration, or any other function performed with finite-state devices, a projection of the output of the transducer that is a projection of the gesture path G is taken. Accordingly, the  
15 projection of the output of the gesture transducer is used to perform the applicable function using finite-state devices.

After the gesture symbols G and the words W are integrated using the finite-state devices G: W and G\_W:M, for example, i.e., after multimodal integration, the gesture path G and  
20 meaning path M in the resulting finite-state device are used to reestablish the connection between the SEM symbols and their specific contents, for example entities or points selected by the user, that are stored in the I path of the gesture transducer I\*G. In particular, in order to reintegrate the specific  
25 contents of the gesture, the gesture transducer I:G is composed with the gesture path G and meaning path M of the device resulting from multimodal integration (I: G o G:M=I: M). In addition, in order to output the meaning, the symbols on the M side are concatenated together. Further, when out-  
30 putting the meaning, if the M symbol is SEM, the symbol on the I side is taken for that arc.

While a three-tape finite-state automaton is feasible in principle, currently available tools for finite-state language processing generally only support two-tape finite-state  
35 automata, i.e., finite-state transducers. Furthermore, speech recognizers typically do not support the use of a three-tape finite-state automaton as a language model. Accordingly, the multimodal recognition and/or meaning system 1000 implements this three-tape finite-state automaton approach by  
40 using a series of finite-state transducers in place of the single three-tape finite-state automaton shown in FIG. 11, as described below. In particular, the three-tape finite-state automaton shown in FIG. 11 and illustrated by the grammar fragment shown in FIG. 12 can be decomposed into an input  
45 component relating the gesture symbols G and the word symbols W and an output component that relates the input component to the meaning symbols M.

FIG. 11 shows a three-tape finite-state automaton that corresponds to the grammar fragment shown in FIG. 12 and that  
50 is usable to recognize the meaning of the various spoken and gestural inputs that can be generated using the exemplary graphical user interfaces displayed using the gesture input portion 410 of the multimodal user input devices 400 shown in FIG. 6. The three-tape finite-state automaton shown in FIG. 11 is decomposed into the gesture-to-speech finite-state  
55 transducer shown in FIG. 13 and the speech/gesture/meaning finite-state transducer shown in FIG. 14.

In various exemplary embodiments of the systems and methods according to this invention, the multimodal gram-  
60 mar contains extensible-markup-language tags. For example, the multimodal grammar shown in FIG. 12 has, for example, the <cmd>, <rest>, <show>, <phone> and <addr> XML tags in the meaning part of the terminals. Therefore, for example, in the finite-state devices used in various exemplary embodi-  
65 ments of the systems and methods according to this invention, some of the transition symbols in the finite-state devices are

extensible-markup-language tags. Functions, such as, for example, multimodal integration, are performed using the finite-state devices. According to the systems and methods according to this invention, if the symbols representing the  
5 output of the performed function are concatenated together, coherent extensible-markup-language expressions can be generated.

The gesture-to-speech finite-state transducer shown in FIG. 13 maps the gesture symbols G to the word symbols W that are expected to coincide with each other. Thus, for  
10 example, in the exemplary embodiment of the multimodal user input device 400 shown in FIG. 6, the verbal pointers “that” and “this” are expected to be accompanied by the deictic gestures made against a department object, an organiza-  
15 tion object or a person object.

The gesture-to-speech transducer shown in FIG. 13 captures the constraints that the gestures, made by the user through the gesture input portion 410 of the multimodal user  
20 input device 400, place on the speech utterance that accompanies those gestures. Thus, a projection of the output tape of the gesture-to-speech finite-state transducer shown in FIG. 13 can be used, in conjunction with the recognized gesture string, such as the recognized gesture string shown in FIG. 16 that represents the gestures illustrated in the exemplary  
25 embodiment of the multimodal user input device 400 shown in FIG. 6, as a language model usable to constrain the possible sequences of words to be recognized by the utterance recognition subsystem 170 of the automatic speech recognition system 100.

The speech/gesture/meaning finite-state transducer shown in FIG. 14 uses the cross-product of the gesture symbols G and the word symbols W as an input component or first tape.  
30 Thus, the gesture-to-speech finite-state transducer shown in FIG. 13 implements the function  $\mathfrak{R}: G \rightarrow W$ . The output or second tape of the speech/gesture/meaning finite-state transducer shown in FIG. 14 contains the meaning symbols M that capture the semantic representation of the multimodal utterance, as shown in FIG. 11 and outlined above. Thus, the  
35 speech/gesture/meaning finite-state transducer shown in FIG. 14 implements the function  $\mathfrak{S}: (G \times W) \rightarrow M$ . That is, the speech/gesture/meaning finite-state transducer shown in FIG. 14 is a finite-state transducer in which gesture symbols and words are on the input tape and the meaning is on the output  
40 tape.

Thus, the gesture-to-speech finite-state transducer and the speech/gesture/meaning finite-state transducers shown in  
45 FIGS. 13 and 14 are used with the speech recognition system 100 and the multimodal parser/meaning recognition system 300 to recognize, parse, and/or extract the meaning from the multimodal inputs received from the gesture and speech input portions 410 and 420 of the multimodal user input device 400.

It should be appreciated that there are any variety of ways in which the multimodal finite-state transducers can be integrated with the automatic speech recognition system 100, the  
50 gesture recognition system 200 and the multimodal parser/meaning recognition system 300. Clearly, for any particular recognition task, the more appropriate approach will depend on the properties of the particular multimodal user input interface 400 through which the multimodal inputs are generated and/or received.

The approach outlined in the following description of FIGS. 11-25 involves recognizing the gesture string first. The recognized gesture string is then used to modify the language  
55 model used by the automatic speech recognition system 100. In general, this will be appropriate when there is limited ambiguity in the recognized gesture string. For example, there will be limited ambiguity in the recognized gesture

string when the majority of gestures are unambiguous deictic pointing gestures. Obviously, if more complex gestures are used, such as the multi-element gestures described in Johnston 1-3, other ways of combining the gesture and speech recognition systems may be more appropriate.

Accordingly, for the specific exemplary embodiment of the multimodal user input device **400** shown in FIG. **6**, the gesture recognition system **200** first processes the incoming gestures to construct a gesture finite-state transducer, such as that shown in FIG. **15**, corresponding to the range of gesture interpretations. As discussed above, the incoming gestures can, for example, be in the form of speech inputs, pen inputs, computer vision inputs, haptic inputs, gaze inputs, and/or body motion inputs.

In particular, in the systems and methods according to this invention, to abstract over specific content of the input gesture, the input gesture is represented as a finite state transducer. This finite-state transducer relates the gesture symbols as well as the specific contents, such as, for example, entities or points on the gesture input portion selected by the user, of the input gesture to the gesture symbols, including a symbol acting as a place holder for the specific contents of the input gesture. More particularly, one side of this gesture finite-state transducer contains the specific contents of the input gesture and the other side of the gesture finite-state transducer contains a symbol, for example, the “SEM” term, which acts as a place holder for the specific content of the input gesture which is on the other side of the gesture finite-state transducer. In addition, in order to perform multimodal integration using finite-state devices, the output projection of the gesture finite-state transducer, i.e., the side without the specific contents of the input gesture, is taken. For example, the output projection of the gesture finite-state transducer shown in FIG. **15** is shown as the gesture finite-state machine in FIG. **16**. As discussed above, after multimodal integration, the gesture finite-state transducer is composed with a projection of the resulting gesture/speech/meaning finite-state machine in order to re-integrate the specific contents of the input gesture.

In addition, in the exemplary embodiments described above with respect to FIG. **6**, the gesture input is unambiguous. Thus, as shown in FIG. **16**, a simple linearly-connected set of states forms the gesture finite-state machine. It should be appreciated that, if the received gestures involved more complex gesture recognition or were otherwise ambiguous, the recognized string of gestures would be represented as a lattice indicating all of the possible gesture recognitions and interpretations for the received gesture stream. Moreover, a weighted finite-state transducer could be used to incorporate the likelihoods of the various paths in such a lattice.

FIG. **17** is a flowchart outlining one exemplary embodiment of a method for combining and converting the various multimodal input streams into a combined finite-state transducer representing the semantic meaning of the combined multimodal input streams according to this invention. Beginning in step **500**, control continues to step **510**, where gesture and speech utterances are input through one or more input devices that together combine to form a multimodal user input device.

Next, in step **515**, a gesture finite-state transducer is generated, from the input gesture utterance. The gesture finite-state transducer relates the specific contents, such as a list of entities, entity identifiers representing the entities selected by the user, or points on the gesture input portion selected by the user, of the gesture and the gesture symbols representing the input gesture to the gesture symbols representing the input gesture, including a symbol acting as a place holder for the specific contents of the input gesture.

This gesture finite-state transducer essentially stores the specific contents of the input gesture, such as entities or points on the gesture input portion, selected by the user, on one side of the gesture finite-state transducer and uses a symbol, such as, for example, the “SEM” term, as a place holder for the specific contents of the input gesture on the other side of gesture finite-state transducer.

Then, in step **520**, a gesture finite-state machine is generated from the gesture finite-state transducer. The gesture finite-state machine is generated by taking a projection of the output of the gesture finite-state transducer. The output of the gesture finite-state transducer contains the gesture symbols including a symbol acting as a place holder for the specific contents of the input gesture. The projection of the output of the gesture finite state transducer, i.e. the gesture finite-state machine, can be used to perform a function, such as multimodal integration, using finite-state devices.

Alternatively, in step **515**, a gesture lattice that represents the input gesture can be generated from the input gesture utterance. The gesture lattice contains gesture symbols and the specific contents of the gesture, such as entities or points selected by the user. In this case, in step **520**, a gesture finite-state transducer is generated from the gesture finite-state machine and the gesture finite-state transducer relates the contents of the gesture lattice to gesture symbols representing the input gesture, including a symbol acting as a place holder for the specific contents of the input gesture.

Next, in step **530**, the gesture finite-state machine is composed with the gesture-to-speech transducer to generate a gesture/language finite-state transducer. For example, in the exemplary embodiment described above, the gesture finite-state machine shown in FIG. **15** is composed with the gesture-to-speech finite-state transducer shown in FIG. **13** to form the gesture/language finite-state transducer shown in FIG. **18**. The gesture/language finite-state transducer represents the relationship between the recognized stream of gestures and all of the possible word sequences that could occur with those gestures of the recognized stream of gestures.

Then, in step **540**, in order to use this information to guide the speech recognition system **100**, a projection of the gesture/language finite-state transducer is generated. In particular, a projection on the output tape or speech portion of the gesture/language finite-state transducer shown in FIG. **18** is taken to yield a finite-state machine.

Next, in step **550**, the speech utterance is recognized using the projection of the gesture/language finite-state transducer as the language model. Using the projection of the gesture/language finite-state transducer as the language model enables the gestural information to directly influence the recognition process performed by the automatic speech recognition system **100**. In particular, as shown in step **560**, the automatic speech recognition system generates a word sequence lattice based on the projection of the gesture/language finite-state transducer in view of the word lattice **155**. In the exemplary embodiment outlined above, using the projection of the gesture/language finite-state transducer shown in FIG. **18** as the language model for the speech recognition process results in the recognized word sequence lattice “phone for these three restaurants”, as shown in FIG. **19**.

Then, in step **570**, the gesture/language finite-state transducer is composed with the recognized word sequences lattice to generate a gesture/speech finite-state transducer. This reintegrates the gesture information that was removed when the projection of the gesture/language finite-state transducer was generated in step **540**. The generated gesture/speech finite-state transducer contains the information both from the speech utterance and the gesture utterance received from the

## 25

various portions of the multimodal user input device 400. For the example outlined above, composing the gesture/language finite-state transducer shown in FIG. 18 with the word sequences lattice shown in FIG. 19 generates the gesture/speech finite-state transducer shown in FIG. 20. Operation then continues to step 580.

Then, in step 580, the gesture/speech finite-state transducer is converted to a gesture/speech finite-state machine. In particular, the gesture/speech finite-state machine combines the input and output tapes of the gesture/speech finite-state transducer onto a single tape. In the exemplary embodiment outlined above, converting the gesture/speech finite-state transducer shown in FIG. 20 results in the gesture/speech finite-state machine shown in FIG. 21.

Next, in step 590, the gesture/speech finite-state machine is composed with the speech/gesture/meaning finite-state transducer shown in FIG. 14 to generate the meaning finite-state transducer shown in FIG. 22. Because the speech/gesture/meaning finite-state transducer relates the speech and gesture symbols to meaning, composing the gesture/speech finite-state machine results in the meaning finite-state transducer which captures the combined semantic meaning or representation contained in the independent modes input using the multimodal user input device. Thus, the meaning of the multimodal input received from the multimodal user input device can be read from the output tape of the meaning finite-state transducer. In the exemplary embodiment outlined above, composing the gesture/speech finite-state machine shown in FIG. 21 with the speech/gesture/meaning finite-state transducer shown in FIG. 14 results in the meaning finite-state transducer shown in FIG. 22. In particular, it should be appreciated that the meaning finite-state transducer shown in FIG. 22 is a linear finite-state transducer that unambiguously yields the meaning. Operation then continues to step 600.

In step 600, as shown in FIG. 23, a projection that factors out speech from the of the meaning finite-state transducer is taken and in step 610, as shown in FIG. 24, the projection is composed with the gesture input finite-state transducer in order to reincorporate the specific contents of the input gesture. Operation then continues to step 620. As shown in FIG. 25, for example, in step 620, a meaning is obtained. In various exemplary embodiments of the systems and methods according to this invention, the string of meaning is a coherent extensible-markup-language expression that can be used by other applications.

The resulting expression can be used by other applications. For example, a graphical user interface that includes a working-city-guide application and navigation screen can, for example, use the resulting markup-language expression to provide users with the requested restaurant and subway information for a city. The user is, for example, able to interact with the graphical interface to, for example, display restaurant listings, restaurant reviews, cuisine/menus, phone numbers, a dynamic map showing locations and street information and/or other information for a restaurant or a set of restaurants.

For example, the string of meaning shown in FIG. 25 when concatenated together results in the following extensible-markup-language expression:

---

```

<cmd>
  <phone>
    <rest> [id1] </rest>
    <rest> [id2, id3] </rest>
  <phone>
</cmd>

```

---

The resulting markup-language expression can be used by, for example, the working-city-guide application, to provide

## 26

the telephone numbers for the selected restaurants. Similar expressions can be used to provide an application program interface (API) to a variety of different applications, for example, such as, for applications related to corporate directory and messaging, applications related to medial information access, applications related to map-based information systems for the military and applications for disaster response and recovery.

It should be appreciated that, in various exemplary embodiments of the systems and methods of this invention, the resulting markup-language expression can be parsed by an applicable parser, i.e., a parser which can parse the markup-language expression. The parsed markup-language expression can be used, for example, by a subsequent application and/or process to further process the information contained therein and/or to execute any commands contained therein.

It should also be appreciated that, in embodiments that use much more complex multimodal interfaces, such as those illustrated in Johnston 1-3, the meaning finite-state transducer may very well be a weighted finite-state transducer having multiple paths between the start and end nodes representing the various possible meanings for the multimodal input and the probability corresponding to each path. In this case, in step 620, the most likely meaning would be selected from the meaning finite-state transducer based on the path through the meaning finite-state transducer having the highest probability. However, it should be appreciated that step 620 is optional and can be omitted. Then, in step 630, the process ends.

As outlined above, the various exemplary embodiments described herein allow spoken language and gesture input streams to be parsed and integrated by a single weighted finite-state device. This single weighted finite-state device provides language models for speech and gesture recognition and composes the meaning content from the speech and gesture input streams into a single semantic representation. Thus, the various systems and methods according to this invention not only address multimodal language recognition, but also encode the semantics as well as the syntax into a single weighted finite-state device. Compared to the previous approaches for integrating multimodal input streams, such as those described in Johnston 1-3, which compose elements from n-best lists of recognition results, the systems and methods according to this invention provide the potential for mutual compensation among the various multimodal input modes.

The systems and methods according to this invention allow the gestural input to dynamically alter the language model used for speech recognition. Additionally, the systems and methods according to this invention reduce the computational complexity of multi-dimensional multimodal parsing.

It should be appreciated that the systems and methods for creating markup-language expressions according to this invention, are not limited to multimodal integration. The systems and methods for creating markup-language expressions according to this invention can be used, for example, to process an input represented by or using a finite-state device in order to generate a markup-language expression. The systems and methods for creating markup-language expressions according to this invention can also, for example, be used to transform one markup-language expression into another markup-language expression.

It should be appreciated that the multimodal recognition and/or meaning system 1000 shown in FIG. 3, and/or each of the gesture recognition system 200, the multimodal parser/meaning recognition system 300 and/or the automatic speech

recognition system **100** can each be implemented on a programmed general purpose computer. However, any or all of these systems can also be implemented on a special purpose computer, a programmed microprocessor or microcontroller and peripheral integrated circuit elements, an ASIC or other integrated circuit, a digital signal processor, a hardwired electronic or a logic circuit such as a discrete element circuit, a programmable logic device such as a PLD, a PLA, a FPGA or a PAL, or the like. In general, any device capable of implementing a finite-state machine that is in turn capable of implementing the flowchart shown in FIG. **15** and/or the various finite-state machines and transducers shown in FIGS. **19-21** and **23-28** can be used to implement one or more of the various systems shown in FIGS. **2-5**.

Thus, it should be understood that each of the various systems and subsystems shown in FIGS. **2-5** can be implemented as portions of a suitably programmed general purpose computer. Alternatively, each of the systems or subsystems shown in FIGS. **2-5** can be implemented as physically distinct hardware circuits within an ASIC, or using a FPGA, a PLD, a PLA, or a PAL, or using discrete logic elements or discrete circuit elements. The particular form each of the systems and/or subsystems shown in FIGS. **2-5** will take is a design choice and will be obvious and predictable to those skilled in the art.

It should also be appreciated that, while the above-outlined description of the various systems and methods according to this invention and the figures focus on speech and gesture as the multimodal inputs, any known or later-developed set of two or more input streams representing different modes of information or communication, such as speech, electronic-ink-based gestures or other haptic modes, keyboard input, inputs generated by observing or sensing human body motions, including hand motions, gaze motions, facial expressions, or other human body motions, or any other known or later-developed method for communicating information, can be combined and used as one of the input streams in the multimodal utterance.

Thus, while this invention has been described in conjunction with the exemplary embodiments outlined above, it is evident that many alternatives, modifications and variations will be apparent to those skilled in the art. Accordingly, the exemplary embodiments of these systems and methods according to this invention, as set forth above, are intended to be illustrative, not limiting. Various changes may be made without departing from the spirit and scope of this invention.

We claim:

**1.** A computing device that generates markup-language-based expressions from input, the computing device comprising:

a processor;

a memory;

a module stored in the memory that controls the processor to receive at least one data input in at least one input mode of a recognition device;

a module stored in the memory that controls the processor to generate at least one input lattice representing the at least one data input;

a module stored in the memory that controls the processor to use at least one finite-state device to perform a function using the at least one generated input lattice, the finite-state device having symbols based on markup-language semantics, wherein performing the function comprises mapping a first markup-language expression to a second markup-language expression using a transducer relating the first-markup-language expression to the second-markup-language expression;

a module stored in the memory that controls the processor to generate a result finite-state device representing a result of the performed function, the result finite-state device relating at least one set of input symbols to at least one set of output symbols;

a module stored in the memory that controls the processor to generate the at least one markup-language expression; and

a module stored in the memory that controls the processor to output the at least one markup-language expression to a processor.

**2.** The computing device of claim **1**, wherein the module stored in the memory that controls the processor to use at least one finite-state device to perform a function using the at least one generated input lattice further uses at least one finite-state device representing a regular grammar containing markup language tags in at least one terminal of the regular grammar.

**3.** The computing device of claim **2**, wherein the module stored in the memory that controls the processor to use at least one finite-state device representing a regular grammar containing markup-language tags in at least one terminal of the regular grammar further uses at least one finite-state device representing a regular grammar containing markup-language tags in a result-part of at least one terminal of the regular grammar.

**4.** The computing device according to claim **2**, wherein:

the module stored in the memory that controls the processor to use at least one finite-state device to perform a function using the at least one generated input lattice further performs multimodal integration using the at least one generated input lattice; and

the module stored in the memory that controls the processor to generate a finite-state device representing a result of the performed function further generates a result-finite-state device relating at least one set of input symbols to at least one set of output symbols representing the multimodal integration.

**5.** The computing device according to claim **2**, wherein the module stored in the memory that controls the processor to use at least one finite-state device to perform a function using the at least one generated input lattice further converts a sequence of words into a sequence of markup-language symbols using a finite-state transducer relating words to markup-language symbols.

**6.** The computing device of claim **1**, wherein the module stored in the memory that controls the processor to receive at least one data input in at least one input mode of a recognition device further inputs at least one of a command and information using at least one of a speech input subsystem, a pen input subsystem, a computer vision input subsystem, a haptic input subsystem, a gaze input subsystem, and a body motion input subsystem.

**7.** The computing device according to claim **1**, wherein the module stored in the memory that controls the processor to use at least one finite-state device to perform a function using the at least one generated input lattice further uses extensible-markup-language-based symbols in the finite state device having symbols based on markup-language semantics.

**8.** The computing device according to claim **1**, further comprising a module stored in the memory that controls the processor to use the at least one generated markup-language expression to perform at least one other function.

**9.** A computing device that generates markup-language-based expressions from an input, the computing device comprising:

a module stored in the memory that controls the processor to receive at least one data input in at least one input mode of a recognition device;

a module stored in the memory that controls the processor to generate at least one input lattice representing the at least one data input;

a module stored in the memory that controls the processor to use at least one finite-state transducer relating a first set of symbols to a second set of symbols to perform a function using the at least one generated input lattice, the second set of symbols having markup-language-based symbols, wherein performing the function further comprises mapping a first-markup-language expression to a second-markup-language expression using a transducer relating the first-markup-language expression to the second-markup-language expression;

a module stored in the memory that controls the processor to generate a result finite-state device representing a result of the performed function, the result finite-state device relating at least one set of input symbol's to at least one set of output symbols;

a module stored in the memory that controls the processor to concatenate output symbols in the at least one set of output symbols of the result finite-state device to generate at least one markup-language expression;

a module stored in the memory that controls the processor to generate the at least one markup-language expression; and

a module stored in the memory that controls the processor to output the at least one generated markup-language expression to a processor.

**10.** The computing device of claim **9**, wherein the module stored in the memory that controls the processor to use at least one finite-state transducer relating a first set of symbols to a second set of symbols further uses at least one finite-state transducer representing a regular grammar containing markup-language tags.

**11.** The computing device of claim **10**, wherein the module stored in the memory that controls the processor to use at least one finite-state transducer representing a regular grammar containing markup-language tags further uses at least one finite-state transducer representing a regular grammar containing markup-language tags in a result-part of at least one terminal of the regular grammar.

**12.** The computing device according to claim **10**, wherein: the module stored in the memory that controls the processor to use at least one finite-state transducer relating a first set of symbols to a second set of symbols further performs multimodal integration using the at least one generated input lattice; and

the module stored in the memory that controls the processor to generate a result-finite-state device representing a result of the performed function further generates a result-finite-state device relating at least one set of input symbols to at least one set of output symbols representing the multimodal integration.

**13.** The computing device of claim **9**, wherein the module stored in the memory that controls the processor to receive at least one data input in at least one input mode of a recognition device further inputs at least one of a command and information using at least one of a speech input subsystem, a pen input subsystem, a computer vision input subsystem, a haptic input subsystem, a gaze input subsystem, and a body motion input subsystem.

**14.** The computing device according to claim **9**, wherein the module stored in the memory that controls the processor to use at least one finite-state transducer relating a first set of

symbols to a second set of symbols to perform a function using the at least one generated input lattice, the at least one finite-state transducer having symbols based on markup-language semantics, further uses extensible-markup-language-based symbols in the at least one finite state transducer having symbols based on markup-language semantics.

**15.** The computing device according to claim **9**, further comprising a module stored in the memory that controls the processor to use the at least one generated markup-language expression to perform at least one other function.

**16.** A computer readable medium storing instructions for controlling a computing device to generate a markup-language-based expression from an input into an unimodal or a multimodal device, the instructions causing the computing device to perform steps comprising:

receiving at least one data input in at least one input mode of a recognition device;

generating at least one input lattice representing the at least one data input;

using at least one finite-state device to perform a function using the at least one generated input lattice, the finite-state device having symbols based on markup-language semantics, wherein performing the function comprises mapping a first markup-language expression to a second markup-language expression using a transducer relating the first-markup-language expression to the second-markup-language expression;

generating a result finite-state device representing a result of the performed function, the result finite-state device relating at least one set of input symbols to at least one set of output symbols;

generating the at least one markup-language expression; and

outputting the at least one markup-language expression to a processor.

**17.** The computer readable medium of claim **16**, wherein using at least one finite-state device to perform a function using the at least one generated input lattice comprises using at least one finite-state device representing a regular grammar containing markup language tags in at least one terminal of the regular grammar.

**18.** The computer readable medium of claim **17**, wherein using at least one finite-state device representing a regular grammar containing markup-language tags in at least one terminal of the regular grammar comprises using at least one finite-state device representing a regular grammar containing markup-language tags in a result-part of at least one terminal of the regular grammar.

**19.** The computer readable medium according to claim **17**, wherein:

using at least one finite-state device to perform a function using the at least one generated input lattice comprises performing multimodal integration using the at least one generated input lattice; and

generating a result-finite-state device representing a result of the performed function comprises generating a result-finite-state device relating at least one set of input symbols to at least one set of output symbols representing the multimodal integration.

**20.** The computer readable medium according to claim **17**, wherein using at least one finite-state device to perform a function using the at least one generated input lattice comprises converting a sequence of words into a sequence of markup-language symbols using a finite-state transducer relating words to markup-language symbols.

**21.** The computer readable medium of claim **16**, wherein receiving at least one data input in at least one input mode of

## 31

a recognition device comprises inputting at least one of a command and information using at least one of a speech input subsystem, a pen input subsystem, a computer vision input subsystem, a haptic input subsystem, a gaze input subsystem, and a body motion input subsystem.

22. The computer readable medium according to claim 16, wherein using at least one finite-state device to perform a function using the at least one generated input lattice comprises using extensible-markup-language-based symbols in the finite state device having symbols based on markup-language semantics.

23. The computer readable medium according to claim 16, wherein the instructions further comprise using the at least one generated markup-language expression to perform at least one other function.

24. A computer readable medium storing instructions for controlling a computing device to generate markup-language-based expressions from an input into an unimodal device or a multimodal device, the instructions comprising:

- receiving at least one data input in at least one input mode of a recognition device;
- generating at least one input lattice representing the at least one data input;

## 32

using at least one finite-state transducer relating a first set of symbols to a second set of symbols to perform a function using the at least one generated input lattice, the second set of symbols having markup-language-based symbols, wherein performing the function further comprises mapping a first-markup-language expression to a second-markup-language expression using a transducer relating the first-markup-language expression to the second-markup-language expression;

generating a result finite-state device representing a result of the performed function, the result finite-state device relating at least one set of input symbol's to at least one set of output symbols;

concatenating output symbols in the at least one set of output symbols of the result finite-state device to generate at least one markup-language expression;

generating the at least one markup-language expression; and

outputting the at least one generated markup-language expression to a processor.

\* \* \* \* \*